arXiv:1609.08879v1 [cs.DS] 28 Sep 2016

Fine-Grained Algorithm Design for Matching

George B. Mertzios[1], André Nichterlein*[1,2], and Rolf Niedermeier[2]

[1]School of Engineering and Computing Sciences, Durham University, UK,
george.mertzios@durham.ac.uk
[2]Institut für Softwaretechnik und Theoretische Informatik, TU Berlin, Germany,
{andre.nichterlein,rolf.niedermeier}@tu-berlin.de

Abstract

Finding maximum-cardinality matchings in undirected graphs is arguably one of the most central graph problems. For m-edge and n-vertex graphs, it is well-known to be solvable in $O(m\sqrt{n})$ time; however, for several applications this running time is still too slow. Improving this worst-case upper bound resisted decades of research. In this paper we mainly focus on parameterizations of the input with respect to several kinds of *distance to triviality*, i.e. how far is the input from some linear-time solvable cases. Our contribution is twofold. First we focus on linear-time fixed-parameter algorithms (with low polynomial parameter dependence). To this end we develop the first linear-time algorithm for maximum matching on *cocomparability* graphs; this algorithm is based on the recently discovered *Lexicographic Depth First Search (LDFS)* and is of independent interest. Using this algorithm we derive an $O(k(n+m))$-time algorithm for general graphs, where k is the vertex deletion distance to cocomparability graphs. Second we focus on *linear-time kernelization*. We start a deeper and systematic study of various "distance to triviality"-parameters for the maximum matching problem. We design linear (and almost linear) time computable kernels of size $O(k)$, $O(k^2)$, $O(k^3)$, and $2^{O(k)}$, respectively, where k is the considered parameter in each case. Investigating linear-time kernelization of a polynomial-time solvable problem, such as maximum matching, leads to a rich number of new and combinatorially interesting challenges. Based on our results, we postulate that maximum matching has the clear potential to become the "drosophila" of "FPT in P studies" analogously to the path-breaking role vertex covering played for classical FPT studies for NP-hard problems.

1 Introduction

"Matching is a powerful piece of algorithmic magic" [33]. The problem Matching is, given a graph, to compute a maximum-cardinality set of nonoverlapping edges. This problem is arguably among the most fundamental graph-algorithmic primitives that are doable in polynomial time. More specifically, on an n-vertex and m-edge graph a maximum matching can be found in $O(m\sqrt{n})$ time [28]. Improving this upper time bound resisted decades of research. Recently, however, Duan and Pettie [10] presented a linear-time algorithm that computes an $(1-\epsilon)$-approximate maximum weight matching (the running time factor depending on ϵ is

*Supported by a postdoc fellowship of the German Academic Exchange Service (DAAD) while at Durham University.

Table 1: Fastest algorithms for MATCHING on special graph classes; $\omega < 2.373$ is the matrix multiplication coefficient.

graph class	running time
general	$O(m\sqrt{n})$ [28], $O(m\sqrt{n}\log(n^2/m)/\log(n))$ [16], $O(n^\omega)$ (randomized) [30]
bipartite	$O(m\sqrt{n})$ [18], $O(n^\omega)$ (randomized) [30, 32], $\tilde{O}(m^{10/7}) = \tilde{O}(m^{1.43})$ [25]
interval	$O(n \log n + m)$ [24]
circular arcs	$O(n \log n)$ [24]
co-interval	$O(n \log n + m)$ [14]
convex bipartite	$O(n)$ [35]
planar	$O(n^{\omega/2})$ (randomized) [29]
strongly chordal	$O(n + m)$ (given the strongly perfect elimination order) [8]
chordal bipartite	$O(n + m)$ [5]
regular	$O(n^2 \log n)$ [37]
cographs	$O(n)$ (given a co-tree) [36]
co-comparability	$O(n + m)$ (Theorem 9 in Section 3)

$\epsilon^{-1}\log(\epsilon^{-1})$). For the unweighted case, it was known that the $O(m\sqrt{n})$ algorithm of Micali and Vazirani [28] implies an $(1-\epsilon)$-approximation running in $O(m\epsilon^{-1})$ time [10]. In this work we take a different route and try to identify relevant special cases that allow for the computation of exact solutions in (almost) linear time. Doing so, however, we focus here on the unweighted case, that is, the computation of maximum-cardinality matchings.

The consideration of more efficiently solvable special cases of finding maximum matchings is not new. For instance, since an augmenting path can be found in linear time [12], the basic augmenting-path based algorithm runs in $O(k(n+m))$ time, where k is the number of edges in the matching. Yuster [37] developed an algorithm with running time $O(rn^2 \log n)$, where r denotes the difference between maximum and minimum vertex degree of the input graph. Moreover, there are linear-time algorithms for computing maximum matchings in special classes of graphs, including convex bipartite graphs [35], strongly chordal graphs [8], and chordal bipartite graphs [5]. We refer to Table 1 for a more thorough overview, also containing results with superlinear running times.

This paper is motivated by the spirit of "parameterization for polynomial-time solvable problems" [15] (also referred to as "FPT in P" or "FPTP" for short). Our contribution is on the identification of faster exact solutions for maximum matchings in relevant special cases. Before describing our specific results, let us remark that Fomin et al. [11] developed an algorithm to compute a maximum matching in graphs of treewidth t (assuming that the corresponding tree decomposition is given) in $O(t^4 n \log n)$ randomized time. We focus on efficient deterministic algorithms whose running time linearly depends on the input size. Moreover, we work toward a more holistic understanding of the *fine-grained* complexity landscape of computing maximum matchings. We have two main types of contributions.

First, we exploit a simple framework which many efficient matching algorithms are based on: Initially compute some matching and then perform augmenting path computations to obtain an optimal solution. As Kaya et al. [20] observed in experimental work with several algorithms for maximum bipartite matching, "elaborate initialization heuristics often have a greater impact than different algorithmic techniques." Instead of relying on heuristics

to compute some initial matching in G, our approach is to use a linear-time algorithm for computing a *maximum* matching in $G - X$ for some vertex subset X. Then, observe that $|X| = k$ augmenting path computations (each doable in linear time) suffice to compute a maximum matching of G in $O(k(n + m))$ time in total. Note that these theoretical findings might also help to better understand and predict the practical success of matching algorithms based on initialization heuristics. This simple strategy motivates the design of (i) linear-time algorithms for maximum matching on special graph classes and (ii) linear-time algorithms for computing (or approximating) the minimum vertex deletion set to specific graph classes.

The second type of results is even stronger rooted in parameterized algorithmics. Here, we follow the paradigm of *kernelization*, that is, provably effective and efficient data reduction. Our aim within FPTP is to devise problem kernels that are computable in (almost) linear time. A particular motivation for this is that with such very efficient kernelization algorithms it is possible to transform *multiplicative* $(f(p)(n + m))$ into *additive* $(f'(p) + O(n + m))$ "(almost) linear-time FPTP" algorithms. Furthermore, the kernelization can be used as a preprocessing to heuristics or approximation algorithms.

Taking a more general point of view, our work can also be interpreted as a serious attempt to establish the faster computation of maximum matchings as the "drosophila" for FPTP studies (which, in a systematic manner, just started) in the same spirit as the computation of minimum vertex covers was in the—meanwhile roughly 25 year-old and very successful—history of parameterized algorithmics[1]. Parameters (acting as secondary measurements beyond mere input size) were key to the flourishing field of studying the fine-grained complexity landscape of NP-hard problems. It is natural to expect similar effects from fine-grained studies for (often, particularly when confronted with big data, only theoretically tractable) polynomial-time solvable problems. Indeed, due to its very nature, we believe that FPTP studies may achieve stronger direct impact on practical computing than standard FPT studies for NP-hard problems so far do. The reason is that *all* aspects of the (polynomial) running time are seriously considered and tuned as much as possible, making aspects such as clever data structures much more relevant.

Our contributions. We present new provably efficient algorithms for Matching (see Table 2 for an overview). In particular we provide new insights into its fine-grained complexity, which is all about getting below the worst case running time $O(m\sqrt{n})$. With our results we aim at establishing Matching as a core problem in the development of the emerging FPTP field of study. Our contributions are twofold.

First we provide in Section 3 a linear-time algorithm for Matching on *cocomparability* graphs. Its relevance with FPTP comes from the framework described above, which uses the iterative augmenting path computations. We use this framework in the spirit of "distance to triviality"-parameterization [4, 17, 22], thus obtaining a linear-time FPTP algorithm; in this case the parameter is the vertex deletion distance to cocomparability graphs.

Apart from its relevance to FPTP, our maximum matching algorithm for cocomparability graphs is also of independent interest. It is a simple greedy algorithm, called *Rightmost Matching (RMM)*, which runs on a specific vertex ordering σ; this ordering is obtained by using (as a preprocessing step) the recently discovered *Lexicographic Depth First Search (LDFS)* algorithm [7]. Interestingly, although the proof of correctness of RMM on cocomparability

[1]Basically every parameterized complexity textbook seriously draws on the rich set of findings concerning vertex covering, these having to a large extent triggered and driven the development of the field.

Table 2: Overview of our results for MATCHING and BIPARTITE MATCHING. Here, k always denotes the corresponding ("distance") parameter.

Parameter	Result
	Results for Bipartite Matching
Distance to chain graphs	$O(k(n+m))$ time algorithm (Section 4) $O(k^3)$-vertex kernel computable in $O(n+m)$ time (Theorem 20)
	Results for Matching
Distance to cocomparability graphs	$O(k(n+m))$ (Theorem 10, requires the distance set to be given in the input)
Feedback edge number	$O(k(n+m))$ algorithm (Theorem 4) kernel of size $O(k)$ computable in $O(n+m)$ time (Theorem 23)
Vertex cover number	$O(k(n+m))$ time algorithm (Theorem 4) kernel of size $O(k^2)$ computable in $O(n+m)$ time (Theorem 25)
Feedback vertex number	$O(k(n+m))$ time algorithm (Theorem 4) kernel of size $2^{O(k)}$ computable in $O(kn)$ time (Theorem 40)

graphs is technically involved, it turns out that RMM computes in a trivial way a maximum matching on interval graphs, when applied on the standard interval graph vertex ordering[2]. Note that the class of interval graphs is a strict subset of the class of cocomparability graphs. A similar phenomenon of extending an interval graph algorithm to cocomparability graphs by using an LDFS preprocessing step has also been observed for the longest path problem [27] and for the minimum path cover problem [6]. Our result on the RMM algorithm, as well as the previous results [6, 27], provide evidence that cocomparability graphs present an "interval graph structure" when they are considered using an LDFS vertex ordering. This may lead to other new and more efficient combinatorial algorithms.

Second we take a deeper look into the issue of kernelization possibilities for MATCHING and we systematically investigate several structural parameterizations. At first we parameterize the input instance with respect to *feedback edge number* (Section 4.2), to *vertex cover number* (Section 4.3), and to *feedback vertex number* (Section 5), which measure the distance to some trivially solvable *sparse* instances (namely an independent set and a forest, respectively). On the one hand, we could relatively easily develop linear-time computable polynomial-size kernels with respect to feedback edge number and vertex cover number. On the other hand, we developed an almost linear-time kernelization algorithm for feedback vertex number, which turned out to be technically very challenging even when only guaranteeing an exponential-size kernel. All these three parameters are among the best-studied and best-understood ones in classical FPT algorithm design for NP-hard problems. However, as it turns out, their investigation in the context of FPTP leads to a rich number of new and combinatorially interesting challenges. This is mainly due to the fact that, in order to achieve a linear (or almost linear) time kernelization algorithm, we often need to use suitable data structures and to carefully design the appropriate data reduction rules to be exhaustively applicable in linear time.

Apart from the above parameters, in Section 4.1 we study BIPARTITE MATCHING with respect to the parameter *distance to chain graphs* (which are a subclass of bipartite graphs).

[2]This is the vertex ordering that results by sorting the intervals according to their left endpoints.

As a maximum matching can be computed in linear time on chain graphs [35], this parameter measures the distance to a trivially solvable *dense* instance (namely a chain graph). In this case, again with quite some technical expenditure, we develop a linear-time computable cubic-size kernel for Bipartite Matching parameterized by the distance to chain graphs.

2 Preliminaries and basic observations

Notation. We use standard notation from graph theory. In particular all paths we consider are simple paths. Two paths P and Q in a graph are called *internally vertex-disjoint* if they are either completely vertex-disjoint or they overlap only in their endpoints. A *matching* in a graph is a set of pairwise disjoint edges. Let $G = (V, E)$ be a graph and let $M \subseteq E$ be a matching in G. A vertex $v \in V$ is called *matched* with respect to M if there is an edge in M containing v, otherwise v is called *free* with respect to M. If the matching M is clear from the context, then we omit "with respect to M". An *alternating path* with respect to M is a path in G such that every second edge of the path is in M. An *augmenting path* is an alternating path whose endpoints are free. It is well known that a matching M is maximum if and only if there is no augmenting path for it. Let $M \subseteq E$ and $M' \subseteq E$ be two matchings in G. We denote by $G(M, M') := (V, M \triangle M')$ the graph containing only the edges in the symmetric difference of M and M', that is, $M \triangle M' := M \cup M' \setminus (M \cap M')$. Observe that every vertex in $G(M, M')$ has degree at most two [18].

A basic observation. For a matching $M \subseteq E$ for G we denote by $M_G^{\max}(M)$ a maximum matching in G with the largest possible overlap with M. That is, $M_G^{\max}(M)$ is a maximum matching in G such that for each maximum matching M' for G it holds that $|M \triangle M'| \geq |M \triangle M_G^{\max}(M)|$. Observe that, if M is a maximum matching for G, then $M_G^{\max}(M) = M$. Furthermore observe that $G(M, M_G^{\max}(M))$ consists of only odd-length paths and isolated vertices, and each of these paths is an augmenting path for M. Moreover the paths in $G(M, M_G^{\max}(M))$ are as short as possible.

Observation 1. *For any path $v_1, v_2, \ldots, v_p$ in $G(M, M_G^{\max}(M))$ it holds that $\{v_{2i-1}, v_{2j}\} \notin E$ for every $1 \leq i < j \leq p/2$.*

Proof. Assume that $\{v_{2i-1}, v_{2j}\} \in E$. Then $v_1, v_2, \ldots, v_{2i-2}, v_{2i-1}, v_{2j}, v_{2j+1}, \ldots, v_p$ is a shorter path which is also an augmenting path for M in G. The corresponding maximum matching M' satisfies $|M \triangle M_G^{\max}(M)| > |M \triangle M'|$, a contradiction to the definition of $M_G^{\max}(M)$. □

A simple standard framework for solving Matching. We describe in the following a generic, augmenting path-based approach to provide FPTP algorithms for Matching. To this end, we use a standard approach to solve Matching: First compute with a fast heuristic a large (but usually suboptimal) matching and then use augmenting path computations to iteratively improve the initial matching to an optimum matching. An augmenting path can be found in linear time (see e. g. Blum [2], Gabow and Tarjan [13], Micali and Vazirani [28]). Thus, the running time of this approach is $O(h + p \cdot (n + m))$ where h is the running time for the initial heuristic and p denotes the number of times the algorithm needs to search for an augmenting path, that is, the size difference of the maximum matching and the initial

matching found by the heuristic. In the following, we want to have $h = O(n + m)$ and we upper-bound p using the following easy observation.

Observation 2. *Let $G = (V, E)$ be a graph with a maximum matching M_G, let $X \subseteq V$ be a vertex subset of size k, and let M_{G-X} be a maximum matching for $G - X$. Then, $|M_{G-X}| \leq |M_G| \leq |M_{G-X}| + k$.*

The correctness of the observation can easily be verified: The part $|M_{G-X}| \leq |M_G|$ is trivial. For the part $|M_G| \leq |M_{G-X}| + k$ observe that M_G is a maximum matching in G and since each vertex appears in at most one edge of M_G, it follows that $|M_G| - k \leq |M_{G-X}|$.

We use Observation 2 in the following way. We look for a large induced subgraph H of the input graph G such that

1. H contains all but k vertices of G for some k and
2. one can compute a maximum matching in H in linear time.

The number k will be the parameter in our FPTP algorithm and it measures the *distance to triviality* [4, 17, 22]. (In our context, "triviality" means an (almost) linear-time solvable case.) Putting these things together, we arrive at the following algorithm:

1. Use the constant-factor linear-time approximation to compute X such that $G - X \in \mathcal{G}$.
2. Compute in linear time a maximum matching M in $G - X$.
3. Start with M as a matching in G and increase the size at most k times to obtain in $O(k \cdot (n + m))$ time a maximum matching for G.

This yields the following central fact:

Observation 3. *Let $\mathcal{G}$ be a graph class. If* Matching *can be solved in linear time on $\mathcal{G}$ and there is a constant-factor linear-time approximation to compute a minimum-size vertex set X such that $G - X \in \mathcal{G}$, then* Matching *can be solved in $O(|X| \cdot (n + m))$ time.*

From the above observation we derive two tasks which we will address in the first part of our work:

1. Find graph classes where Matching is (almost) linear-time solvable.
2. Provide (almost) linear-time (constant-factor approximation) algorithms to compute the vertex deletion distance to these graph classes.

Combining known results from the literature with Observation 3 yields the following:

Theorem 4. Matching *can be solved in $O(k(n + m))$ time, where k is one of the following parameters: feedback vertex number, feedback edge number, vertex cover number.*

Proof. Computing a maximum matching in a tree is easily doable in linear time. Bar-Yehuda et al. [1] provided a linear-time constant-factor approximation for computing a minimum feedback vertex set. Thus, Observation 3 implies an $O(k(n + m))$ time algorithm for Matching parameterized by the feedback vertex number k. Observe that in any graph both the feedback edge number and the vertex cover number are lower bounded by the feedback vertex number. Furthermore, a minimum feedback edge set can be computed in linear time by computing

a minimum spanning tree (all edges not in the minimum spanning tree form the feedback edge set). Finally, there is a simple linear-time factor-two approximation for computing a minimum vertex cover: pick an arbitrary uncovered edge e, add both endpoints of e to the vertex cover, delete the endpoints, and repeat until no edge is left in the graph. Thus, we arrive at $O(k(n+m))$ time algorithms for MATCHING parameterized by the feedback edge number and by the vertex cover number. □

3 A linear-time algorithm for cocomparability graphs

In this section we present a linear-time greedy algorithm that computes a maximum matching on interval graphs (a subclass of comparability graphs) and on cocomparability graphs, see Theorem 9. A graph $G = (V, E)$ is an *interval* graph if we can assign to each vertex of G a closed interval on the real line such that two vertices are adjacent in G if and only if the corresponding two intervals intersect. A *comparability* graph is a graph whose edges can be transitively oriented, i.e. if $u \to v$ and $v \to w$ then $u \to w$. A *cocomparability* graph G is a graph whose complement $\overline{G}$ is a comparability graph. The class of interval graphs is strictly included in the class of cocomparability graphs [3]. Intuitively, we can transitively orient the "non-edges" of an interval graph using the ordering of non-intersecting intervals from left to right, as follows. Consider three intervals I_a, I_b, I_c in an interval representation of an interval graph. If I_a lies completely to the left of I_b, and I_b lies completely to the left of I_c, then also I_a lies completely to the left of I_c.

To begin with, we sketch a simple greedy linear-time algorithm for computing a maximum matching M on interval graphs. First, given an interval graph G with n vertices and m edges, we compute in $O(n+m)$ time an interval representation of G and, at the same time, we also sort the intervals according to their left endpoint [31]. Our algorithm works as follows. Initially $M = \emptyset$ and all vertices are marked as "unvisited". At every iteration we visit the unvisited vertex (interval) x which has the rightmost left endpoint, and we mark x as "visited". Then, if x does not have any unvisited neighbor in G, then we proceed by visiting the next unvisited vertex (interval) with the currently rightmost left endpoint; note that this vertex is now different from x, as x has been already marked as visited. Otherwise, if x has at least one unvisited neighbor in G, then we proceed by visiting the unvisited neighbor y of x with the rightmost left endpoint; in this case we mark y as "visited" and we also add the edge $\{x, y\}$ to the computed matching M.

We call the above algorithm *Rightmost-Matching (RMM)*. It can be executed in $O(n+m)$ time; with a simple exchange argument we can show that the matching M returned by RMM is indeed maximum in G. This algorithm implicitly uses the following vertex ordering that characterizes interval graphs. This ordering can also result by sorting (in linear time [31]) the intervals according to their left endpoints.

Lemma 5 ([31]). *$G = (V, E)$ is an interval graph if and only if there exists a vertex ordering σ of G (called an* I-ordering*) such that, for all $x <_\sigma y <_\sigma z$, if $\{x, z\} \in E$, then also $\{x, y\} \in E$.*

As our main result of this section, we prove below that this simple RMM algorithm can actually compute a maximum matching in the much greater class of cocomparability graphs, once it is applied to a specific vertex ordering σ; see Algorithm 1 for a description of RMM in the more general case where the input is a cocomparability graph. This ordering σ is obtained by applying (as a preprocessing step) the recently discovered *Lexicographic Depth*

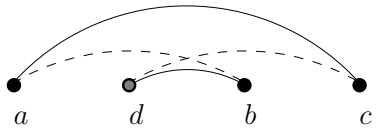

Figure 1: A good triple (a, b, c) and its d-vertex, in the vertex ordering $\sigma = (a, d, b, c)$. The edges $\{a, c\}$ and $\{d, b\}$ are indicated with solid lines and the non-edges $\{a, b\}$ and $\{d, c\}$ with dashed lines.

First Search (LDFS) algorithm [7] to a vertex ordering π that characterizes cocomparability graphs (see Definition 1 and Lemma 6). LDFS is a variation of the well known Depth First Search (DFS) algorithm, which appropriately assigns labels to the vertices and uses the lexicographic order over these labels as a tie-breaking rule.

Before we proceed with our algorithm and its analysis, we first present in Definition 1 the *umbrella-free* ordering that characterizes the more general class of cocomparability graphs. This ordering is a direct generalization of the I-ordering for interval graphs, see Lemma 5. It is worth noting here that, although there exists a linear-time algorithm to *compute* an umbrella-free ordering π of a given cocomparability graph [26], the fastest known algorithm to *verify* that a given vertex ordering is indeed umbrella-free needs the same time as boolean matrix multiplication (see [34] for a discussion of this issue).

Definition 1 ([23])**.** Let $G = (V, E)$ be a graph. An ordering π of the vertices V is an *umbrella-free* ordering (or a *CO-ordering*) if for all $x <_\pi y <_\pi z$ it holds that if $\{x, z\} \in E$, then $\{x, y\} \in E$ or $\{y, z\} \in E$ (or both).

Lemma 6 ([23])**.** *A graph $G = (V, E)$ is a cocomparability graph if and only if there exists an umbrella-free ordering π of V.*

As a first step for our algorithm we compute from the given umbrella-free ordering π a new vertex ordering σ by applying a preprocessing step to π which is based on *LDFS* algorithm [7]. We present in Definition 3 below the notion of an *LDFS-ordering*. It has been proved that a vertex ordering σ is an LDFS-ordering if and only if σ can be returned by an application of the LDFS algorithm [7]. Before we present Definition 3, which is crucial for our analysis, we first present the next auxiliary definition.

Definition 2 ([7])**.** Let $G = (V, E)$ be a graph and σ an ordering of V. Let $a, b, c \in V$ be three vertices such that $a <_\sigma b <_\sigma c$, $\{a, c\} \in E$, and $\{a, b\} \notin E$. If there exists a vertex d such that $a <_\sigma d <_\sigma b$, $\{d, b\} \in E$, $\{d, c\} \notin E$, then (a, b, c) is a *good triple*, otherwise it is a *bad triple*. Furthermore, for a good triple (a, b, c), this vertex d is called its *d-vertex*.

Definition 3 ([7])**.** Let $G = (V, E)$ be a graph. An ordering σ of V is an *LDFS ordering* if σ has no bad triple.

An example of a good triple (a, b, c) and a d-vertex of it is depicted in Figure 1. Note that the input to the standard LDFS algorithm is a graph G and a starting vertex v. As an additional tie-breaking rule (usually referred to as a "+ rule"), LDFS can use a given vertex ordering π. In this case, LDFS$^+$ visits at every step the *rightmost* unvisited vertex in π among all vertices that have the currently *lexicographically largest* label. The resulting ordering σ is then denoted $\sigma =$LDFS$^+(G, \pi)$; see [6, 7, 27] for a more detailed discussion. Given a vertex ordering π of a graph G with n vertices and m edges, LDFS$^+(G, \pi)$ can be

Algorithm 1: RMM(G, σ).

Input: A cocomparability graph $G = (V, E)$ and an LDFS umbrella-free ordering σ of V.
Output: A vertex ordering $\widehat{\sigma}$ and a maximum matching M of G

```
1 Label all vertices "unvisited"; i ← 0; M ← ∅
2 while there are unvisited vertices do
3     Pick the rightmost unvisited vertex x in σ and mark x as "visited"
4     i ← i + 1; σ̂(i) ← x                          // add vertex x to the ordering σ̂
5     if x has at least one unvisited neighbor then
6         Pick the rightmost unvisited neighbor y of x and mark y as "visited"
7         i ← i + 1; σ̂(i) ← y                      // add vertex y to the ordering σ̂
8         M ← M ∪ {{x, y}}                                          // match x and y
9 return the ordering σ̂ and the matching M.
```

computed in $O(n + m)$ time [21]. In the preprocessing step for our algorithm we compute the ordering $\sigma = \text{LDFS}^+(G, \pi)$ where π is the given umbrella-free ordering of G. Then, this LDFS vertex ordering σ remains umbrella-free [6].

Once we have computed the LDFS umbrella-free ordering $\sigma = \text{LDFS}^+(G, \pi)$, we apply our simple linear-time algorithm *Rightmost-Matching (RMM)*, see Algorithm 1, to compute a new vertex ordering $\widehat{\sigma}$ and a maximum matching M of G. RMM is a simple "greedy" algorithm which operates as follows. At every step it visits the rightmost unvisited vertex x in σ and it marks x as visited. Then, if x does not have any unvisited neighbor, then RMM proceeds at the next step by visiting again the currently unvisited vertex in σ; note that this vertex is now different from x, as x has been already marked as visited. Otherwise, if x has at least one unvisited neighbor, then RMM visits after x its rightmost unvisited neighbor y in the ordering σ and it also adds the edge $\{x, y\}$ to the computed matching M. Since any I-ordering of an interval graph is also an LDFS umbrella-free ordering (see Lemma 5 and Definitions 1 to 3), note that Algorithm 1 also works with an interval graph G and an I-ordering σ of G as input; this is actually exactly the same RMM algorithm as we sketched above for interval graphs.

In the remainder of this section, we show that the matching M returned by RMM(G, σ) is indeed a maximum matching of G. The proof is by contradiction and uses an appropriate *potential function* f that is defined over all matchings of G, see Definition 4. In particular, we prove in Theorem 9 that M is a maximum matching of G.

Definition 4 (potential function)**.** Let $G = (V, E)$ be a cocomparability graph and σ be an LDFS umbrella-free ordering of $V = \{v_1, \ldots, v_n\}$ with $v_1 <_\sigma \ldots <_\sigma v_n$. Let M be a matching of G. We define the potential function $f(M) := \sum_{i=1}^{n} g(v_i)$, where for each $v_i \in V$:

$$g(v_i) := \begin{cases} 0, & \text{if } \{v_i, v_j\} \in M \text{ and } i < j, \\ (i - j) \cdot (n + 1)^i, & \text{if } \{v_i, v_j\} \in M \text{ and } j < i, \\ i \cdot (n + 1)^i, & \text{if } v_i \text{ is not matched within } M. \end{cases}$$

Note by Definition 4 that, for the empty matching, we have $f(\emptyset) = \sum_{i=1}^{n} i \cdot (n + 1)^i$. Then, as we add an edge $\{\{v_i, v_j\}\}$ to the current matching M, where $j < i$, we have that $f(M \cup \{\{v_i, v_j\}\}) = f(M) - j \cdot (n + 1)^i$. The next observation illustrates the connection between the RMM algorithm and this potential function.

Observation 7. *The matching M returned by RMM(G, σ) minimizes the function $f(M)$.*

Proof. Let M' be a matching such that $f(M')$ is minimum. Consider the rightmost vertex v_n in the ordering σ and let v_i be the rightmost neighbor of v_n in σ. Assume that $\{v_i, v_n\} \notin M'$. Then let M'' be the matching obtained by removing from M' the edges that v_i and v_n are matched to, and by adding to it the edge $\{v_i, v_n\}$. Then it is easy to see that $f(M'') < f(M')$ which is a contradiction. Thus $\{v_i, v_n\} \in M'$. We can now recursively apply the same argument in the induced subgraph $G \setminus \{v_i\} \setminus \{v_n\}$, which eventually implies that M' is the matching returned by RMM(G, σ). □

Before we prove the main result of this section in Theorem 9, we need to prove a crucial technical lemma (Lemma 8). In the next definition we introduce for every vertex v the induced subgraph $G_\sigma(v)$ with respect to the ordering σ, which is fundamental for the statement and the proof of Lemma 8.

Definition 5. Let $G = (V, E)$ be a cocomparability graph and $\sigma = (v_1, v_2, \ldots, v_n)$ be an LDFS umbrella-free vertex ordering of G. For every $v_i \in V$ the graph $G_\sigma(v_i)$ is the induced subgraph of G on the vertices $\{v_1, v_2, \ldots, v_i\}$.

Lemma 8. *Let $G = (V, E)$ be a cocomparability graph and σ be an LDFS umbrella-free ordering of V. Let M be a maximum matching of G such that $f(M)$ is minimum among all maximum matchings. Then, there is no quadruple (a, b, c, x) of vertices in G satisfying all of the following conditions:*

1. *$a <_\sigma b <_\sigma c \leq_\sigma x$,*
2. *$\{a, c\}, \{b, c\} \in E$ and $\{a, b\} \notin b$,*
3. *$\{a, c\} \in M$,*
4. *there is no odd-length alternating path from a to b within $G_\sigma(x)$,*
5. *there is no odd-length alternating path from a to any free vertex v within $G_\sigma(x)$, and*
6. *there is no odd-length alternating path from b to any free vertex v within $G_\sigma(x)$.*

Proof. Let G, σ, and M as described in the statement of the lemma. Towards a contradiction let (a, b, c, x) be a quadruple of vertices satisfying all six conditions of the lemma. Fix now vertex x. Among all such quadruples with fixed x, let (a, b, c, x) be such that a is leftmost in σ, that is, for any other such quadruple (a', b', c', x) we have $a \leq_\sigma a'$. Since σ is an LDFS ordering, it follows from Conditions 1 and 2 and Definitions 2 and 3 that there is a vertex d such that $a <_\sigma d <_\sigma b$, $\{d, b\} \in E$, and $\{d, c\} \notin E$. Since $\{d, c\} \notin E$ and σ is umbrella-free, it follows that $\{a, d\} \in E$. Observe that d is matched in M as otherwise Conditions 5 and 6 would be violated. Thus, there is a vertex $e \in V$ with $\{e, d\} \in M$. Now we distinguish three cases with respect to the position of e in the ordering σ.

Case 1: $c <_\sigma e$. In this case we have that $a <_\sigma c <_\sigma e$, $\{d, e\} \in E$, and $\{d, c\} \notin E$. Thus, since σ is umbrella-free, it follows that $\{c, e\} \in E$. However, in this case for the matching $M' = (M \setminus \{\{a, c\}, \{d, e\}\}) \cup \{\{e, c\}, \{a, d\}\}$ we have that $f(M') < f(M)$, a contradiction.

Case 2: $a <_\sigma e <_\sigma c$. If $\{a, e\} \in E$, then there exists the alternating path (a, e, d, b) of length 3 from a to b within $G_\sigma(x)$, which is a contradiction to Condition 4. Thus $\{a, e\} \notin E$.

Furthermore $\{c, e\} \in E$, since σ is umbrella-free and $\{a, c\} \in E$. Thus for the matching $M' = (M \setminus \{\{a, c\}, \{d, e\}\}) \cup \{\{e, c\}, \{a, d\}\}$ we have that $f(M') < f(M)$, a contradiction.

Case 3: $e <_\sigma a$**.** In this case it follows similarly to Case 2 that $\{a, e\} \notin E$ (proof by contradiction due to Condition 4). Furthermore observe that $\{e, d\}, \{a, d\} \in E$ and $\{e, d\} \in M$. Thus the triple (e, a, d) satisfies Conditions 1 to 3. Furthermore, if there exists an odd-length alternating path from e to a within $G_\sigma(x)$, then this alternating path can be extended through d to an odd-length alternating path from e to b within $G_\sigma(x)$, which is a contradiction to Condition 4. Therefore there is no odd-length alternating path from e to a within $G_\sigma(x)$. Similarly, odd-length alternating paths from e (resp. from a) to a free vertex v within $G_\sigma(x)$ are excluded as well due to Condition 5 (resp. due to Condition 6). Thus the quadruple (e, a, d, x) satisfies the six conditions of the lemma and it holds that $e <_\sigma a$, a contradiction to the choice of the initial quadruple (a, b, c, x). □

Theorem 9. *Given a cocomparability graph G, Algorithm 1 returns a maximum matching M of G in $O(n + m)$ time.*

Proof. Let $G = (V, E)$ be a cocomparability graph, and let σ be an umbrella-free LDFS ordering of G. First we prove that Algorithm 1 runs in $O(n+m)$ time. During the execution of the algorithm we maintain the *unvisited vertices* in an array A (initially of size n), according to their position in σ. Furthermore we maintain for each vertex u its *unvisited neighbors* in an array N_u (initially of size $\deg(u)$), again according to their position in σ. Once we have computed the ordering σ, the construction of array A can be done in $O(n)$ time. The construction of all arrays N_u, where $u \in V$, can be done in $O(n + m)$ time as follows. We initialize $N_u = \emptyset$ for every $u \in V$. Then we iterate for each vertex $u \in V$ in array A from left to right. For every such vertex u we scan (in an arbitrary order) through its neighborhood $N(u)$, and for each $v \in N(u)$ we append vertex u in the array N_v.

Algorithm 1 can be clearly executed in $O(n)$ time. The rightmost unvisited vertex x in Algorithm 1 can be found in $O(1)$ time as the rightmost vertex in the array A. Once x is detected in Algorithm 1, x is removed from A also in $O(1)$ time. Furthermore x is removed from all arrays N_u, where $\{x, u\} \in E$, in $O(\deg(x)$ time. Moreover, Algorithm 1 can be clearly executed in $O(1)$ time. The if-condition of Algorithm 1 can be checked in $O(1)$ time by just checking whether the array N_x is empty. Similarly to Algorithm 1, Algorithm 1 can be executed in $O(\deg(y))$ time. Furthermore, each of Algorithms 1 to 1 can be clearly executed in $O(1)$ time. Summarizing, the total running time of Algorithm 1 is $O(n + \sum_{u \in V} \deg(u)) = O(n + m)$ time.

For the correctness part, the proof is done by contradiction. Let M be the matching returned by RMM(G, σ). Assume towards a contradiction that M is not a maximum-size matching. For the rest of the proof, let M' denote a maximum matching that minimizes $f(M')$ among all maximum matchings of G. Let x be the rightmost vertex in σ on which M' differs from M'. Then x is matched in at least one of the two matchings M and M'. Now we distinguish three cases with respect to the vertex that is matched with x in M and M'.

Case 1: x is matched in M' to some $y \in V$ but is free in M. Then M and M' also differ at vertex y. Thus $y <_\sigma x$, since x is the rightmost vertex in which M and M' differ. Consider the iteration t of Algorithm 1 during which the algorithm visits x. Suppose that vertex y is matched in M with a vertex z at an earlier iteration $t' < t$. Then M differs from M' also at vertex z. If $z <_\sigma x$, then Algorithm 1 visits x at an earlier iteration than z, which is a contradiction to the assumption on z. Therefore $x <_\sigma z$. This is a contradiction to the assumption that x is the rightmost vertex in σ in which M differs from M'.

Case 2: x is matched in M to some vertex $y \in V$ but is free in M'. If y is free in M', then the matching $M' \cup \{\{x, y\}\}$ is greater than M', which is a contradiction to the maximality assumption on M'. Therefore y is matched in M' to some vertex $z \in V$. Note that M and M' differ also on y and z. Thus, it follows by the choice of x that $y <_\sigma x$ and $z <_\sigma x$. Consider now the matching $M'' := (M' \setminus \{\{y, z\}\}) \cup \{\{x, y\}\}$, which is maximum since $|M''| = |M'|$. However $f(M'') < f(M')$, which is a contradiction to the assumption on the minimality of $f(M')$.

Case 3: $\{x, y\} \in M$ and $\{x, z\} \in M'$ with $z \neq y$. Then M and M' differ also on y and z. Thus, it follows by the choice of x that $y <_\sigma x$ and $z <_\sigma x$. Consider the iteration t of Algorithm 1 during which the algorithm visits x. Suppose that vertex z is matched in M with a vertex p at an earlier iteration $t' < t$. Then M differs from M' also at vertex p. If $p <_\sigma x$, then Algorithm 1 visits x at an earlier iteration than p, which is a contradiction to the assumption on p. If $x <_\sigma p$, then we have again a contradiction to the assumption that x is the rightmost vertex in σ in which M differs from M'. Thus z is unmatched in M at the iteration t of Algorithm 1 during which the algorithm visits x. Furthermore z is also unvisited at iteration t, since $z <_\sigma x$. Now, if $y <_\sigma z$, then Algorithm 1 would not match x to y at the execution of Algorithm 1, which is a contradiction. Therefore $z <_\sigma y$.

Suppose that y is free in M'. Then $M'' := (M' \setminus \{\{x, z\}\}) \cup \{\{x, y\}\}$ is another maximum matching, for which $f(M'') < f(M')$, which is a contradiction to the assumption on M'. Therefore y is matched in M' to some vertex $w \in V$. If $\{w, z\} \in E$ then the matching $M'' := (M' \setminus \{\{x, z\}, \{w, y\}) \cup \{\{x, y\}, \{z, w\}$ is another maximum matching, for which $f(M'') < f(M')$, which a contradiction to the choice of M'. Therefore $\{z, w\} \notin E$.

Suppose that there exists within $G_\sigma(x)$ (Definition 5) an odd-length alternating path P_0 with respect to M' from w to z. Then, swapping all edges on the path P_0, removing $\{x, z\}$ and $\{w, y\}$ from M', and adding $\{x, y\}$ yields another maximum matching M''. Recall that x is the rightmost vertex in which M and M' differ. Thus, since the alternating path P_0 belongs to the induced subgraph $G_\sigma(x)$, it follows that $f(M'') < f(M')$, which a contradiction to the choice of M'. Thus there exists within $G_\sigma(x)$ no odd-length alternating path with respect to M' from w to z.

Similarly, suppose that there exists within $G_\sigma(x)$ an odd-length alternating path P_1 with respect to M' from w (resp. from z) to a free vertex v. Then, swapping all the edges on the path P_1, removing $\{x, z\}$ and $\{w, y\}$ from M' and adding $\{x, y\}$ yields another maximum matching M'' such that $f(M'') < f(M')$, which is again a contradiction to the choice of M'. Thus there exists within $G_\sigma(x)$ no odd-length alternating path with respect to M' from w (resp. from z) to a free vertex v.

Now suppose that $w <_\sigma z$. That is, $w <_\sigma z <_\sigma y$, where $\{w, y\} \in E$ and $\{z, w\} \notin E$. Therefore $\{z, y\} \in E$, since σ is umbrella-free. Thus, since $\{w, y\} \in M'$, it follows that the quadruple (w, z, y, x) satisfies all six conditions in the statement of Lemma 8. This is a contradiction to Lemma 8, since M' is assumed to be a maximum matching of G such that $f(M')$ is minimum among all maximum matchings.

Finally suppose that $z <_\sigma w$. Recall that M differs from M' in w, since $\{y, w\} \in M'$ and $\{x, y\} \in M$. Thus, since x is the rightmost vertex in σ in which M differs from M', it follows that $w <_\sigma x$. That is, $z <_\sigma w <_\sigma x$, where $\{x, z\} \in E$ and $\{z, w\} \notin E$. Therefore $\{w, x\} \in E$, since σ is umbrella-free. Thus, since $\{x, z\} \in M'$, it follows that the quadruple (z, w, x, x) satisfies all six conditions in the statement of Lemma 8. This is again a contradiction to Lemma 8, since M' is assumed to be a maximum matching of G such that $f(M')$ is minimum among all maximum matchings.

Summarizing, the matching M returned by Algorithm 1 is a maximum matching. □

The next theorem follows now directly by Observation 3 and Theorem 9.

Theorem 10. MATCHING *can be solved in* $O(k \cdot (n+m))$ *time given a vertex set subset* X *of size* k *such that deleting* X *yields a cocomparability graphs.*

4 Linear-time polynomial-size kernels

In this and the next section, we investigate the possibilty of efficient and effective preprocessing for MATCHING and BIPARTITE MATCHING. To this end, we use the well-established notion of kernelization from parameterized complexity as described below.

A *parameterized problem* is a set of instances (I, k) where $I \in \Sigma^*$ for a finite alphabet Σ, and $k \in \Sigma^*$ is the *parameter*. We say that two instances (I, k) and (I', k') of parameterized problems P and P' are *equivalent* if (I, k) is a yes-instance for P if and only if (I', k') is a yes-instance for P'. A *kernelization* is an algorithm that, given an instance (I, k) of a parameterized problem P, computes in polynomial time an equivalent instance (I', k') of P (the *kernel*) such that $|I'| + k' \leq f(k)$ for some computable function f only depending on k. We say that f measures the *size* of the kernel, and if $f(k) \in k^{O(1)}$, we say that P admits a polynomial kernel. Often, a kernel is achieved by applying polynomial-time executable data reduction rules. We call a data reduction rule $\mathcal{R}$ *correct* if the new instance (I', k') that results from applying $\mathcal{R}$ to (I, k) is equivalent to (I, k). An instance is called *reduced* with respect to some data reduction rule if further application of this rule has no effect on the instance. As kernelization is usually defined for decision problems, we use in the remainder of the paper the decision version of MATCHING. That is, we are given as input a graph G and an integer s, and the task is to decide whether a matching of size at least s exists in G.

Since solving the given instance and returning a trivial yes- or no-instance always produces a constant size kernel, we are looking for kernelization algorithms that are *faster* than the algorithms solving the problem. For NP-hard problems, each kernelization algorithm, since running in polynomial time, is (presumably) faster than any solution algorithm. This is, of course, no longer true when applying kernelization to a polynomial-time solvable problem like MATCHING. While the focus of classical kernelization for NP-hard problems is mostly on improving the *size* of the kernel, we particularly emphasize that for polynomially solvable problems it now becomes crucial to also focus on the *running time* of the kernelization algorithm. More specifically, the multivariate analysis framework should also be applied to the kernelization algorithm itself. For example, a kernelization algorithm running in $O(k^5 n)$ time might be preferable to another one running in $O(n^3)$ time. This is an important distinguishing feature from the kernelization of NP-hard problems. In this paper we present kernelization algorithms for MATCHING which run in *linear time* (this section) or in *almost linear time* (i.e. in $O(kn)$ time where k is the parameter, see next section).

We remark that in classical parameterized complexity for NP-hard problems it is known that a problem is fixed-parameter tractable if and only if it admits a kernel [9]. We do not know whether a analogous statement holds for the polynomial setting. More specifically, we do not know whether e. g. the $O(k(n+m))$ algorithm for MATCHING with respect to the parameter "distance k to cocomparability graphs" (see Theorem 10) implies any kernel that can be computed faster than directly computing a maximum matching.

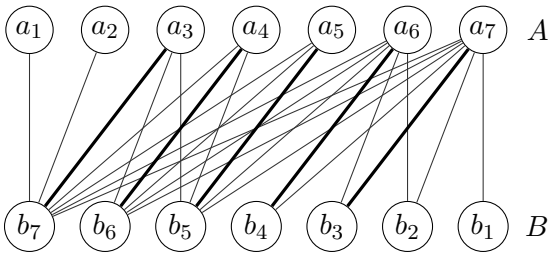

Figure 2: A chain graph. Note that the ordering of the vertices in A is going from left to right while the ordering of the vertices in B is going from right to left. The reason for these two orderings being drawn in different directions is that a maximum matching can be drawn as parallel edges, see e. g. the bold edges. In fact, Algorithm 2 computes such matchings with the matched edges being parallel to each other.

Recall that for each of the parameters our simple framework implies an algorithm with running time $O(k(n+m))$. For feedback edge number and vertex cover number, this follows from Theorem 4. For the third parameter, this follows from the facts that (i) chain graphs are a subclass of convex graphs, (ii) Matching is solvable on the later in linear time [35], and (iii) we can approximate the corresponding vertex deletion set in linear time (see Lemma 12 below).

4.1 Parameter distance to chain graphs

In this section, we show a cubic vertex-kernel for Bipartite Matching parameterized by the "distance k to chain graphs". To this end, we first define chain graphs.

Definition 6 ([3])**.** Let $G = (A, B, E)$ be a bipartite graph. Then G is a *chain graph* if each of its two color classes A, B admits a linear order w.r.t. neighborhood inclusion, that is, $A = \{a_1, \ldots, a_\alpha\}$ and $B = \{b_1, \ldots, b_\beta\}$ where $N(a_i) \subseteq N(a_j)$ and $N(b_i) \subseteq N(b_j)$ whenever $i < j$.

Observe that if the graph G contains twins, then there is more than one linear order w.r.t. neighborhood inclusion. To avoid ambiguities, we fix for the vertices of the color class A (resp. B) in a chain graph $G = (A, B, E)$ one linear order $\prec_A$ (resp. $\prec_B$) such that, for two vertices $u, v \in A$ (resp. $u, v \in B$), if $u \prec_A v$ (resp. if $u \prec_B v$) then $N(u) \subseteq N(v)$. In the remainder of the section we consider a bipartite representation of a given chain graph $G = (A, B, E)$ where the vertices of A (resp. B) are ordered according to $\prec_A$ (resp. $\prec_B$) from left to right (resp. from right to left), as illustrated in Figure 2. For simplicity of notation we use in the following $\prec$ to denote the orderings $\prec_A$ and $\prec_B$ whenever the color class is clear from the context.

We next show, that we can approximate the parameter set in linear time.

Lemma 11 ([3])**.** *A bipartite graph is a chain graph if and only if it does not contain an induced* $2K_2$.

Lemma 12. *There is a linear-time factor-4 approximation for the problem of deleting a minimum number of vertices in a bipartite graph in order to obtain chain graph.*

Algorithm 2: An algorithm that computes a maximum matching M in the chain graph G such that all edges in M are parallel (see Fig. 2 for a visualization.)

Input: A chain graph $G = (V, E)$, $V = A \cup B$, $A = \{a_1, \ldots, a_\alpha\}$ and $B = \{b_1, \ldots, b_\beta\}$ with $N(a_i) \subseteq N(a_j)$ and $N(b_i) \subseteq N(b_j)$ for $i < j$.
Output: A maximum matching of G where all matched edges are parallel.

1 Compute the size s of a maximum matching in G using an algorithm of Steiner and Yeomans [35]
2 $M \leftarrow \{\{a_{\alpha-s+1}, b_\beta\}, \{a_{\alpha-s+2}, b_{\beta-1}\}, \ldots, \{a_\alpha, b_{\beta-s+1}\}\}$
3 **return** M.

Proof. Let $G = (A, B, E)$ be a bipartite graph. We compute a set $S \subseteq A \cup B$ such that $G - S$ is a chain graph and S is at most four times larger than a minimum size of such a set. The algorithm iteratively tries to find a $2K_2$ and deletes the four corresponding vertices until no further $2K_2$ is found. Since in each $2K_2$, by Lemma 11, at least one vertex needs to be removed, the algorithm yields the claimed factor-4 approximation.

The details of the algorithm are as follows: First, it initializes $S = \emptyset$ and sorts the vertices in A and in B by their degree; the vertices in $A = \{a_1, \ldots, a_\alpha\}$ in increasing order and the vertices $B = \{b_1, \ldots, b_\beta\}$ in decreasing order, that is, $\deg(a_1) \leq \ldots \leq \deg(a_\alpha)$ and $\deg(b_1) \leq \ldots \leq \deg(b_\beta)$. Since the degree of each vertex is at most $\max\{\alpha, \beta\}$, this can be done in linear time with e. g. Bucket Sort. At any stage the algorithm deletes all vertices of degree zero and all vertices which are adjacent to all vertices in the other partition. The deleted vertices are not added to S since these vertices can not participate in a $2K_2$. Next, the algorithm recursively processes the vertices in A in a nondecreasing order of their degrees. Let $a \in A$ be a minimum-degree vertex and let $b \in B$ be a neighbor of a. Since b is not adjacent to all vertices in A (otherwise b would be deleted), there is a vertex $a' \in A$ that is not adjacent to b. Since $\deg(a) \leq \deg(a')$ it follows that a' has a neighbor b' that is not adjacent to a. Hence, the four vertices a, a', b, b' induce only two edges: $\{a, b\}$ and $\{a', b'\}$ and thus form a $2K_2$. Thus, the algorithm adds the four vertices to S, deletes them from the graph, and continues with a vertex in A that has minimum degree.

As to the running time, we now show that, after the initial sorting, the algorithm considers each edge only twice: Selecting a and b as described above can be done in $O(1)$ time. To select a', the algorithm simply iterates over all vertices in A until it finds a vertex that is not adjacent to b. In this way at most $\deg(b) + 1$ vertices are considered. Similarly, by iterating over the neighbors of a', one finds b'. Hence, the edges incident to a, a', b, and b' are used once to find the vertices and a second time when these vertices are deleted. Thus, using appropriate data structures, the algorithm runs in $O(n + m)$ time. □

In the rest of this subsection, we provide a linear-time computable kernel for Bipartite Matching with respect to the parameter vertex deletion distance to chain graphs. Let $G = (A, B, E)$ be the bipartite input graph, where $V = A \cup B$, and let $X \subseteq V$ be a vertex subset such that $G - X$ is a chain graph. By Lemma 12, we can compute an approximate X in linear time. The kernelization algorithm is as follows: First, compute a specific maximum matching $M_{G-X} \subseteq E$ in $G - X$ with Algorithm 2 where all edges in M_{G-X} are "parallel" and all matched vertices are consecutive in the ordering $\prec_A$ and $\prec_B$, see also Fig. 2. Since in convex graphs matching is linear-time solvable [35] and convex graphs are a super class of chain graphs, this can be done in $O(n + m)$ time. We use M_{G-X} in our kernelization

algorithm to obtain some local information about possible augmenting paths. For example, each augmenting path has at least one endpoint in X. Forming this into a data reduction rule, with s denoting the size of a maximum matching, yields the following.

Reduction Rule 4.1. *If $|M_{G-X}| \geq s$, then return a trivial yes-instance; if $s > |M_{G-X}| + k$, then return a trivial no-instance.*

The correctness of the above reduction rule follows from Observation 2.

Without loss of generality, we will be able to assume that each vertex in X is either matched in M with another vertex in X or with a "small-degree vertex" in $G-X$. This means that an augmenting path starting at some vertex in X will "enter" the chain graph $G-X$ in a small-degree vertex. We now formalize this concept. For a vertex $x \in X$ we define $N^{V\setminus X}_{\text{small}}(x)$ to be the set of the k neighbors of x in $V \setminus X$ with the smallest degree, formally,

$$N^{V\setminus X}_{\text{small}}(x) := \{w \in N(x) \setminus X \mid k > |\{u \in N(x) \setminus X \mid u \prec w\}|\}.$$

Lemma 13. *Let $G = (V, E)$ be a bipartite graph and let $X \subseteq V$ be a vertex set such that $G-X$ is a chain graph. Then, there exists a maximum matching M_G for G such that every matched vertex $x \in X$ is matched to a vertex in $N^{V\setminus X}_{\text{small}}(x) \cup X$.*

Proof. Assume, towards a contradiction, that there is no such matching M_G. Let M'_G be a maximum matching for G that maximizes the number of vertices $x \in X$ that are matched to a vertex in $N^{V\setminus X}_{\text{small}}(x) \cup X$, that is, let M'_G maximize $|\{x \in X \mid \{u, x\} \in M'_G \wedge u \in N^{V\setminus X}_{\text{small}}(x) \cup X\}|$. Let $x \in X$ be a vertex that is not matched with any vertex in $N^{V\setminus X}_{\text{small}}(x) \cup X$, that is, x is matched to a vertex $u \in V \setminus (N^{V\setminus X}_{\text{small}}(x) \cup X)$. If there is an unmatched vertex $w \in N^{V\setminus X}_{\text{small}}(x)$ in M'_G, then the matching $M''_G := M'_G \cup \{\{x, w\}\} \setminus \{\{u, x\}\}$ is a maximum matching with more vertices $x \in X$ (compared to M'_G) that are matched to a vertex in $N^{V\setminus X}_{\text{small}}(x) \cup X$, a contradiction. Hence, assume that there is no free vertex in $N^{V\setminus X}_{\text{small}}(x)$. Since $|N^{V\setminus X}_{\text{small}}(x)| = |X| = k$, it follows that at least one vertex $w \in N^{V\setminus X}_{\text{small}}(x)$ is matched to a vertex $v \in V \setminus X$. Observe that, by definition of $N^{V\setminus X}_{\text{small}}(x)$, we have $N_{G-X}(w) \subseteq N_{G-X}(u)$. Thus, we have $\{u, v\} \in E$ and thus, $M''_G := M'_G \cup \{\{x, w\}, \{u, v\}\} \setminus \{\{u, x\}, \{w, v\}\}$ is a maximum matching with more vertices in X (compared to M'_G) fulfilling the condition of the lemma, a contradiction. □

Based on Lemma 13, we can provide our next data-reduction rule.

Reduction Rule 4.2. *Let (G, s) be an instance reduced with respect to Reduction Rule 4.1 and let $x \in X$. Then delete all edges between x and $V \setminus N^{V\setminus X}_{\text{small}}(x)$.*

Clearly, Reduction Rule 4.2 can be exhaustively applied in $O(n+m)$ time by one iteration over A and B in the order $\prec$.

The high-level idea of our kernel is as follows: Keep in the kernel all vertices of X. For each vertex $x \in X$ keep all vertices in $N^{V\setminus X}_{\text{small}}(x)$ and if a kept vertex is matched, then keep also the vertex with whom it is matched. Denote with K the set of the vertices kept so far in $V \setminus X$.

Consider an augmenting path $P = x, a_1, b_1, \ldots, a_\ell, b_\ell, y$ from a vertex $x \in B \cap X$ to a vertex $y \in A \cap X$. Observe that if $a_1 \prec a_\ell$, then also $\{b_1, a_\ell\} \in E$ and thus $P' = x, a_1, b_1, a_\ell, b_\ell, y$ is an augmenting path. Furthermore, the vertices in the augmenting path P'

are a subset of $K \cup X$ and, thus, by keeping these vertices (and the edges between them), we also keep the augmenting path P' in our kernel. Hence, it remains to consider the more complicated case that $a_\ell \prec a_1$. To this end, we next show that in certain "areas" of the chain graph $G - X$ the number of augmenting paths "passing through" such an area is bounded.

Definition 7. Let $G = (A, B, E)$ be a chain graph and let M be a matching in G. Furthermore let $a \in A, b \in B$ with $\{a, b\} \in M$. Then $\#\mathrm{lmv}(b, M)$ (resp. $\#\mathrm{rmv}(a, M)$) is the number of neighbors of b (resp. of a) that are to the left of a (resp. to the right of b); formally,

$$\#\mathrm{lmv}(b, M) := |\{a' \in N(b) \mid a' \prec a\}|, \qquad \#\mathrm{rmv}(a, M) := |\{b' \in N(a) \mid b' \prec b\}|.$$

In Definition 7 the terms "left" and "right" refer to the ordering of the vertices of A and B in the bipartite representation of G, as illustrated in Figure 2. The abbreviation #rmv (#lmv) stands for "number of vertices **r**ight (**l**eft) of the **m**atched **v**ertex". We set $\#\mathrm{lmv}(a, M) := \#\mathrm{lmv}(b, M)$ and $\#\mathrm{rmv}(b, M) := \#\mathrm{rmv}(a, M)$. Finally, we define $\#\mathrm{rmv}(a_1, a_2, M) := \min_{a_1 \prec a' \prec a_2}\{\#\mathrm{rmv}(a', M)\}$ for $a_1, a_2 \in A$ and $\#\mathrm{lmv}(b_1, b_2, M) := \min_{b_2 \prec b' \prec b_1}\{\#\mathrm{lmv}(b', M)\}$ for $b_1, b_2 \in B$.

Lemma 14. *Let $G = (A, B, E)$ be a chain graph and M be a maximum matching for G computed by Algorithm 2. Let $a, b \in V$ with $\{a, b\} \in M$. Then the number of vertex-disjoint alternating paths that (1) start and end with edges not in M and that (2) have endpoints left of a and right of b is at most* $\min\{\#\mathrm{lmv}(b, M), \#\mathrm{rmv}(a, M)\}$.

Proof. We prove the case $\#\mathrm{lmv}(b, M) \leq \#\mathrm{rmv}(a, M)$, that is, $\min\{\#\mathrm{lmv}(b, M), \#\mathrm{rmv}(a, M)\} = \#\mathrm{lmv}(b, M)$. The case $\#\mathrm{lmv}(b, M) > \#\mathrm{rmv}(a, M)$ follows by symmetry (with switched roles of a and b). Let #aug denote the number of vertex-disjoint alternating paths from $\{a' \in A \mid a' \prec a\}$ to $\{b' \in B \mid b' \prec b\}$ such that the first and last edge are not in M. Furthermore, let $a^b_1, \ldots, a^b_{\#\mathrm{lmv}(b,M)}$ be the neighbors of b that are to the left of a, that is, $a^b_1 \prec a^b_2 \prec \ldots \prec a^b_{\#\mathrm{lmv}(b,M)} \prec a$. Since G is a chain graph it follows that no vertex $a' \in A$ with $a' \prec a^b_1$ is adjacent to any vertex $b' \in B$ with $b' \preceq b$. Furthermore, for any edge $\{a', b'\} \in E$ with $a \prec a'$ and $b \prec b'$ it follows from the construction of M (see Algorithm 2) that $\{a', b'\} \notin M$. Hence, any of these alternating paths has to contain at least one vertex from $a^b_1, \ldots, a^b_{\#\mathrm{lmv}(b,M)}$. Since the alternating paths are vertex-disjoint it follows that $\#\mathrm{aug} \leq \#\mathrm{lmv}(b, M)$. □

From the previous lemma, we directly obtain the following.

Lemma 15. *Let $G = (A, B, E)$ be a chain graph and let M be the maximum matching for G computed by Algorithm 2. Let $a_1, a_2 \in A$ and $b_1, b_2 \in B$ with $\{a_1, b_1\}, \{a_2, b_2\} \in M$ with $a_1 \prec a_2$. Then there are at most $\#\mathrm{lmv}(b_1, b_2, M)$ vertex-disjoint alternating paths that (1) start and end with edges not in M and that (2) have endpoints left of a_1 and right of b_2.*

Lemma 15 states that the number of augmenting paths passing through the area between a_1 and a_2 is bounded. Using this, we want to replace this area by a gadget with $O(k)$ vertices. To this end, we need further notation. For each kept vertex $v \in K$, we may also keep some vertices to the right and to the left of v. We call these k vertices the *left buffer* (*right buffer*) of v.

Definition 8. Let $G = (A, B, E)$ be a chain graph and let M be the maximum matching for G computed by Algorithm 2. Let $a_1, a_2 \in A$ and $b_1, b_2 \in B$ with $\{a_1, b_1\}, \{a_2, b_2\} \in M$ and $a_1 \prec a_2$. Then the (at most) $\#\mathrm{lmv}(b_1, b_2, M)$ vertices to the right of a_1 form the *right buffer* $B^r(a_1, M)$ of a_1; formally,

$$B^r(a_1, M) := \{a \in A \mid a_1 \prec a \wedge |\{a' \in A \mid a_1 \prec a' \prec a\}| \leq \min\{\#\mathrm{lmv}(b_1, b_2, M), k\}\}.$$

Analogously,

$$B^\ell(a_2, M) := \{a \in A \mid a \prec a_2 \wedge |\{a' \in A \mid a \prec a' \prec a_2\}| \leq \min\{\#\mathrm{lmv}(b_1, b_2, M), k\}\},$$
$$B^r(b_1, M) := \{b \in B \mid b \prec b_1 \wedge |\{b' \in B \mid b \prec b' \prec b_1\}| \leq \min\{\#\mathrm{lmv}(b_1, b_2, M), k\}\},$$
$$B^\ell(b_2, M) := \{b \in B \mid b_2 \prec b \wedge |\{b' \in B \mid b_2 \prec b' \prec b\}| \leq \min\{\#\mathrm{lmv}(b_1, b_2, M), k\}\}.$$

Note that in Definition 8 each of the sets $B^r(a_1, M)$, $B^\ell(a_2, M)$, $B^r(b_1, M)$, and $B^\ell(b_2, M)$ depends on all four vertices a_1, a_2, b_1, b_2; we omit these dependencies from the names for the sake of presentation.

Reduction Rule 4.3. *Let (G, s) be an instance reduced with respect to Reduction Rule 4.1. Let $a_1, a_2 \in K \cap A$ with $a_1 \prec a_2$ and $\{a_1, b_1\}, \{a_2, b_2\} \in M_{G-X}$ such that $A' := \{a \in A \mid a_1 \prec a \prec a_2\}$ is of size at least $2\min\{\#\mathrm{lmv}(b_1, b_2, M), k\} + 1$ and $A' \cap K = \emptyset$. Then delete all vertices in $A' \setminus (B^r(a_1, M_{G-X}) \cup B^\ell(a_2, M_{G-X}))$ and their matched neighbors in B, add all edges between the vertices in the right buffer of a_1 and the vertices in the left buffer of b_2, and decrease s by the number of removed matched vertex pairs.*

Lemma 16. *Reduction Rule 4.3 is correct and can be exhaustively applied in $O(n+m)$ time.*

Proof. We first introduce some notation and provide some general observations. Let a_1, a_2, b_1, and b_2 be as stated in Reduction Rule 4.3. Denote by A' (resp. B') the set of vertices between a_1 and a_2 (resp. between b_1 and b_2). Further denote by $A'_D \subseteq A'$ and $B'_D \subseteq B'$ the sets of deleted vertices. Note that $|A'| = |B'|$ and $|A'_D| = |B'_D|$ since M_{G-X} was produced by Algorithm 2. Denote the vertices in the buffers of a_1, a_2, b_1, and b_2 by $B^x(y_z, M_{G-X}) = \{y^x_1, \ldots, y^x_{\min\{\#\mathrm{lmv}(b_1,b_2,M),k\}}\}$ for $x \in \{r, \ell\}, y \in \{a, b\}, z \in [2]$, and $x = r \iff z = 1$.

Since the input instance is reduced with respect to Reduction Rule 4.1, it follows that $s - k \leq |M_{G-X}| < s$. Denote by $M_{G'-X} := M_{G-X} \cap E'$ the matching obtained from M_{G-X} by deleting all edges not in the reduced graph G'. Recall that s is reduced by the number of matched edges that are removed. We next show in Claims 17 and 18 that the input instance $(G = (V, E), s)$ is a yes-instance if and only if the produced instance $(G' = (V', E'), s')$ is a yes-instance. Before we present these two claims, observe that there is a perfect matching between the vertices in A'_D and B'_D, and thus

$$s - |M_{G-X}| = s' - |M_{G'-X}|. \tag{1}$$

Claim 17. *If (G, s) is a yes-instance, then (G', s') is a yes-instance.*

Proof. Recall that $M^{\max}_G(M_{G-X})$ for G minimizes the size of $M_{G-X} \triangle M^{\max}_G(M_{G-X})$. Since (G, s) is a yes-instance it holds that $|M^{\max}_G(M_{G-X})| \geq s$. For brevity we set $G^M := G(M_{G-X}, M^{\max}_G(M_{G-X})) = (V, M_{G-X} \triangle M^{\max}_G(M_{G-X}))$. Note that G^M is a graph that only contains odd-length paths. We will show that there are as many vertex-disjoint augmenting

paths for $M_{G'-X}$ in G' as there are paths in G^M. This will show that G' contains a matching of size

$$|M_{G'-X}|+|M_G^{\max}(M_{G-X})|-|M_{G-X}| \geq |M_{G'-X}|+s-|M_{G-X}| \overset{(1)}{=} |M_{G'-X}|+s'-|M_{G'-X}| = s'.$$

To this end, observe that all paths that do not use vertices in $V \setminus V' = A'_D \cup B'_D$ are also contained in G'. Thus, consider the paths in G^M that use vertices in $V \setminus V'$. Denote by $\mathcal{P}^M$ the set of all paths in G^M using vertices in $V \setminus V'$ and set $t := |\mathcal{P}^M|$. Consider now an arbitrary $i \in [t]$, and let $P_i^M \in \mathcal{P}^M$. Denote by $v_1^i, v_2^i, \ldots, v_{p_i}^i$ the vertices in P_i^M in the corresponding order, that is, v_1^i and $v_{p_i}^i$ are the endpoints of P_i^M and we have $\{v_{2j}^i, v_{2j+1}^i\} \in M_{G-X}$ for all $j \in [p_i/2]$. Observe that exactly one endpoint of P_i^M is in A and the other endpoint is in B, since P_i^M is an odd-length path. Assume without loss of generality that $v_1^i \in A$ and $v_{p_i}^i \in B$. Thus, the vertices in P_i^M with odd (even) index are in A (B).

We next show that for any two vertices v_j^i, v_ℓ^i of P_i^M with $j < \ell < p_i - 1$ and both being in $A \setminus X$, it follows that $v_j^i \prec v_\ell^i$. First, observe that if $j = 1$, then $v_j^i \in A \setminus X$ is a free vertex wrt. M_{G-X}. Since v_ℓ^i is matched wrt. M_{G-X} and since M_{G-X} is computed by Algorithm 2, it follows that $v_1^i \prec v_\ell^i$. Thus, assume that $j > 1$ and $\ell > 1$ (thus $j \geq 3$ and $j > 1$). Assume toward a contradiction that $v_\ell^i \prec v_j^i$. Since $\ell < p_i - 1$, we have $v_{\ell+1}^i \in B \setminus X$ and since $G - X$ is a chain graph, it follows that $\{v_j^i, v_{\ell+1}^i\} \in E$, a contradiction to Observation 1. Thus, $v_j^i \prec v_\ell^i$.

We next show that the path P_i^M contains at least one vertex v_j^i left of a_1 and at least one vertex v_ℓ^i right of b_2. Recall that M_{G-X} was computed by Algorithm 2 and, thus, the free vertices are the smallest wrt. the ordering $\prec$ (see also Fig. 2). Thus, if one endpoint of P_i^M is in $(A \cup B) \setminus X$, then this vertex is either v_1^i and left of a_1 or it is $v_{p_1}^i$ and right of b_2. Thus, assume that the endpoints of P_i^M are in X. We showed in the previous paragraph that $v_3^i \prec v_5^i \prec \ldots \prec v_{p_1-3}^i$. Thus, we also have $v_{p_i-2}^i \prec v_{p_i-4}^i \prec \ldots \prec v_2^i$ since M_{G-X} is computed by Algorithm 2. Since we assumed that some vertices of P_i^M to be in $V \setminus V'$, it follows that for at least one vertex v_{2j+1}^i it holds that $a_1 \prec v_{2j+1}^i \prec a_2$. Furthermore since by assumption no vertex between a_1 and a_2 or between b_1 and b_2 is in K, it follows that $v_3^i \prec a_1$ (since $v_3^i \in K$) and $v_{p_i-2}^i \prec b_2$ (since $v_{p_i-2}^i \in K$).

For each $i \in [t]$ denote by $a^{P_i^M}$ the last vertex on the path P_i^M that is not right of a_1, that is, $a^{P_i^M}$ is the vertex on P_i^M such that for each vertex $a' \in A \setminus X$ that is in P_i^M it holds that $a_1 \prec a'$ or $a' \preceq a^{P_i^M}$. It follows from the previous paragraph that $a^{P_i^M}$ exists. Analogously to $a^{P_i^M}$, for each $i \in [t]$ denote by $b^{P_i^M}$ the first vertex on the path P_i^M that is not left of b_2, that is, $b^{P_i^M}$ is the vertex on P_i^M such that for each vertex $b' \in B$ that is in P_i^M it holds that $b_2 \prec b'$ or $b' \preceq b^{P_i^M}$. This means that in G^M there is for each $i \in [t]$ an alternating path from $a^{P_i^M}$ to $b^{P_i^M}$ starting and ending with non-matched edges and all these paths are pairwise vertex-disjoint. We show that also in G' there are pairwise vertex-disjoint alternating paths from $a^{P_i^M}$ to $b^{P_i^M}$. Assume without loss of generality that $a^{P_1^M} \prec a^{P_2^M} \prec \ldots \prec a^{P_t^M}$. Since in each path P_i^M, $i \in [t]$, the successor of $a^{P_i^M}$ is to the right of b_1, it follows that $a^{P_i^M}$ has at least i neighbors right of b_1. Since the right buffer of b_1 contains the $\#\mathrm{lmv}(b_1, b_2, M_{G-X}) \geq t$ (see Lemma 15) vertices to the right of b_1, we have $\{a^{P_i^M}, b_i^r\} \in E$. By symmetry, we have $\{b^{P_i^M}, a_i^\ell\} \in E$. Recall that M_{G-X} forms a perfect matching between $B^r(b_1, M_{G-X})$ and $B^r(a_1, M_{G-X})$ as well as between $B^\ell(a_2, M_{G-X})$ and $B^\ell(b_2, M_{G-X})$. Since Reduction Rule 4.3 added all edges between $B^r(a_1, M_{G-X})$ and $B^\ell(b_2, M_{G-X})$ to E', it follows that each path P_i^M can be completed as follows: $a^{P_i^M}, b_i^r, a_i^r, b_i^\ell, a_i^\ell, b^{P_i^M}$; note that exactly the

edges $\{b_i^r, a_i^r\}$ and $\{b_i^\ell, a_i^\ell\}$ are in M_{G-X}. Thus, each path in $\mathcal{P}^M$ can be replaced by an augmenting path for $M_{G'-X}$ in G' and all these augmenting paths are vertex-disjoint. Thus, there are as many augmenting paths for $M_{G'-x}$ in G' as there are paths in G^M and therefore (G', s') is a yes-instance. □

Claim 18. *If (G', s') is a yes-instance, then (G, s) is a yes-instance.*

Proof. Let $M_{G'}$ a maximum matching for G'. Observe that $|M_{G'}| \geq s'$. We construct a matching M_G for G as follows. First, copy all edges from $M_{G'} \cap E$ into M_G. Second, add all edges from $M_{G-X} \cap \binom{V \setminus V'}{2}$, that is, a perfect matching between A'_D and B'_D is added to M_G. Observe that if all edges in $M_{G'}$ are also in E, then M_G is a matching of size s in G. Thus, assume that some edges in $M_{G'}$ are not in E, that is, $\{a_{i_1}^r, b_{j_1}^\ell\}, \ldots, \{a_{i_t}^r, b_{j_t}^\ell\} \in M_{G'} \setminus E$ for some $t \in [\#\mathrm{lmv}(b_1, b_2, M_{G-X})]$. Observe that $s - |M_G| \leq t$. Clearly, the vertices $a_{i_1}^r, \ldots, a_{i_t}^r, b_{j_1}^r, \ldots, b_{j_t}^r$ are free with respect to M_G.

We show that there are t pairwise vertex-disjoint augmenting paths from $\{a_{i_1}^r, \ldots, a_{i_t}^r\}$ to $\{b_{j_1}^\ell, \ldots, b_{j_t}^\ell\}$; note, however, that these paths are not necessarily from $a_{i_r}^r$ to $b_{j_r}^\ell$, where $r \in [t]$. To this end, recall that by definition of $\#\mathrm{lmv}(b_1, b_2, M_{G-X})$ each vertex $b \in B$ with $b_2 \prec b \prec b_1$ has at least $\#\mathrm{lmv}(b_1, b_2, M_{G-X})$ neighbors to the left of its matched neighbor. This allows us to iteratively find augmenting paths as follows: To create the q^{th} augmenting path P_q start with some vertex $b_{j_q}^\ell$. Denote by v the last vertex added to P_q (in the beginning we have $v = b_{j_q}^\ell$). If $v \in A$, then add to P_q the neighbor matched to v. If $v \in B$, then do the following: if v is adjacent to a vertex $a \in \{a_{i_1}^r, \ldots, a_{i_t}^r\}$, then add a to P_q, otherwise add the leftmost neighbor of v to P_q. Repeat this process until P_q contains a vertex from $\{a_{i_1}^r, \ldots, a_{i_t}^r\}$. After we found P_q remove all vertices of P_q from G. If $q < t$, then continue with P_{q+1}. Observe that any two vertices of P_q that are in A have at least $\#\mathrm{lmv}(b_1, b_2, M_{G-X}) - 1$ other vertices of A between them (in the ordering of the vertices of A, see Figure 2). Thus, after a finite number of steps, P_q will reach a vertex in $\{a_{i_1}^r, \ldots, a_{i_t}^r\}$. Furthermore, it follows that after removing the vertices of P_q it holds that $\#\mathrm{lmv}(b_1, b_2, M_{G-X})$ is decreased by exactly one: P_q contains for each vertex $b \in B$ at most one vertex among the $\#\mathrm{lmv}(b_1, b_2, M_{G-X})$ neighbors of b that are directly to the left of its matched neighbor in M_{G-X}. Thus, in each iteration we have $\#\mathrm{lmv}(b_1, b_2, M_{G-X}) > 0$. It follows that the above procedure constructs t vertex-disjoint augmenting paths from $\{a_{i_1}^r, \ldots, a_{i_t}^r\}$ to $\{b_{j_1}^\ell, \ldots, b_{j_t}^\ell\}$. Hence, G contains a matching of size s and thus (G, s) is a yes-instance. □

The correctness of the data reduction rule follows from the previous two claims. It remains to prove the running time. To this end, observe that the matching M_{G-X} is given. Computing all degrees of G can be done in $O(n+m)$ time. Also $\#\mathrm{lmv}(v, M_{G-X})$ can be computed in linear time: For each vertex $b \in B$ one has to check for each neighbor of b whether it is to the left of b's matched neighbor and to adjust $\#\mathrm{lmv}(b, M_{G-X})$ accordingly. Furthermore, computing $\#\mathrm{lmv}(b_1, b_2, M_{G-X})$ and removing the vertices in A'_D and B'_D can be done in $O(\sum_{b \in B'} \deg(b) + \sum_{a \in A'} \deg(a))$ time. Thus, Reduction Rule 4.3 can be exhaustively applied in $O(n + m)$ time. □

We next upper-bound the number of free vertices with respect to M_{G-X}. Let

$$A_{\text{free}} := \{a \in A \mid a \text{ is free with respect to } M_{G-X}\}$$

and

$$A_{\text{free}}^k := \{a \in A_{\text{free}} \mid |\{a' \in A_{\text{free}} \mid a \preceq a'\}| \leq k\},$$

that is, A^k_{free} contains the k rightmost free vertices in $A \setminus X$. Observe that all vertices in A^k_{free} are left of M_{G-X}. Analogously, denote by B^k_{free} the set containing the k leftmost free vertices in $B \setminus X$.

Reduction Rule 4.4. *Let (G, s) be an instance reduced with respect to Reduction Rule 4.1. Then delete all vertices in $\left(A_{free} \setminus (K \cup A^k_{free})\right) \cup \left(B_{free} \setminus (K \cup B^k_{free})\right)$.*

Lemma 19. *Reduction Rule 4.4 is correct an can be applied in $O(n + m)$ time.*

Proof. The running time is clear. It remains to show the correctness. Let (G, s) be the input instance reduced with respect to Reduction Rule 4.1 and let (G', s) be the instance produced by Reduction Rule 4.4. We show that deleting the vertices in $A_{\text{free}} \setminus (K \cup A^k_{\text{free}})$ yields an equivalent instance. It then follows from symmetry that deleting the vertices in $B_{\text{free}} \setminus (K \cup B^k_{\text{free}})$ yields also an equivalent instance.

We first show that if (G, s) is a yes-instance, then also the produced instance (G', s) is a yes-instance. Let (G, s) be a yes-instance and M_G be a maximum matching for G. Clearly, $|M_G| \geq s$. Observe that for each removed vertex $a \in A_{\text{free}} \setminus (K \cup A^k_{\text{free}})$ it holds that every vertex $a' \in A^k_{\text{free}}$ is to the right of a, that is, $a \prec a'$ and thus $N_{G-X}(a) \subseteq N_{G-X}(a')$. Since (G, s) is reduced with respect to Reduction Rule 4.1, it follows that $|M_{G-X}| \geq |M_G| - k$. Thus, there exist at most k augmenting paths for M_{G-X} in G. If none of these augmenting paths ends in a vertex $a \in A_{\text{free}} \setminus (K \cup A^k_{\text{free}})$, then all augmenting paths exist also in G' and thus (G', s) is a yes-instance. If one of these augmenting paths, say P, ends in a, then at least one vertex $a' \in A^k_{\text{free}}$ is not endpoint of any of these augmenting paths. Since $a \notin K$ it follows from Lemma 13 that the neighbor b of a on P is indeed in $B \setminus X$. Since $N_{G-X}(a) \subseteq N_{G-X}(a')$, it follows that $\{a', b\} \in E$ and thus we can replace a by a' in the augmenting path. By exhaustively applying the above exchange argument, it follows that we can assume that none of the augmenting paths uses a vertex in $A_{\text{free}} \setminus (K \cup A^k_{\text{free}})$. Thus, all augmenting paths are also contained in G' and hence the resulting instance (G', s) is still a yes-instance.

Finally observe that if (G', s) is a yes-instance, then also (G, s) is a yes-instance, since G' is a subgraph of G. Thus any matching of size s in G' is also a matching in G. Hence (G, s) is a yes-instance. □

Theorem 20. Matching *on bipartite graphs admits a cubic-vertex kernel with respect to the vertex deletion distance to chain graphs. The kernel can be computed in linear time.*

Proof. Let (G, s) be the input instance with $G = (V, E)$, the two partitions $V = A \cup B$, and $X \subseteq V$ such that $G - X$ is a chain graph. If X is not given explicitly, then use the linear-time factor-four approximation provided in Lemma 12 to compute X. The kernelization is as follows: First, compute M_{G-X} in linear time with Algorithm 2. Next compute the set K. Then, apply Reduction Rules 4.1, 4.3 and 4.4. By Lemmas 16 and 19, this can be done in linear time. Let b^K_ℓ the leftmost vertex in $K \cap B$ and a^K_r rightmost vertex in $A \cap K$. Let a^K_ℓ and b^K_r be their matched neighbors. Since $|K| \leq k^2$ and we reduced the instance with respect to Reduction Rule 4.3, it follows that the number of vertices between a^K_ℓ and a^K_r as well as the number of vertices between b^K_ℓ and b^K_r is at most $2k^3$, respectively. Furthermore, there are at most $2k$ free vertices left in $V \setminus X$ since we reduced the instance with respect to Reduction Rule 4.4. It remains to uper-bound the number of matched vertices left of b^K_ℓ and right of a^K_r.

Observe that all vertices left of b^K_ℓ are matched with respect to M_{G-X}. If there are more than $2k$ vertices to the left of b^K_ℓ, then do the following: Add four vertices $a_\ell, b_\ell, x^a_\ell, x^b_\ell$ to V.

The idea is that $\{a_\ell, b_\ell\}$ should be an edge in M_{G-X} such that $a_\ell \in A$ and $b_\ell \in B$ are in K and there is no vertex left of b_ℓ. This means we add these vertices to simulate the situation where the leftmost vertex in $B \setminus X$ is also in K. To ensure that a_ℓ and b_ℓ are in K and that they are not matched with some vertices in G, we add x_ℓ^a and x_ℓ^b to X and make x_ℓ^a respectively x_ℓ^b to their sole neighbors. In this way, we ensure that there is maximum matching in the new graph that is exactly two edges larger than the maximum matching in the old graph. In this new graph we can then apply Reduction Rule 4.3 to reduce the number of vertices between b_ℓ and b_ℓ^K. Formally, we add the following edges. Add $\{a_\ell, x_\ell^a\}, \{b_\ell, x_\ell^b\}$ to E. Add all edges between b_ℓ and the vertices in $B \setminus X$. Let a be the rightmost vertex in A_{free}^k. Then, add edges between a_ℓ and $N_{G-X}(a)$. Set $b' \prec b_\ell$ for each $b' \in B \setminus X$, set $a \prec a_\ell$ for each vertex $a \in A_{\text{free}}$, and set $a_\ell \prec a'$ for each matched vertex in $A \setminus X$. Furthermore, add $\{a_\ell, b_\ell\}$ to M_{G-X} and add a_ℓ and b_ℓ to K. Finally, increase s by two. Next, apply Reduction Rule 4.3 in linear time, then remove $a_\ell, b_\ell, x_\ell^a, x_\ell^b$ and reduce s by two. After this procedure, it follows that there are at most $2k$ vertices left of b_ℓ^K. If there are more than $2k$ vertices right of the rightmost vertex a_r^K in $A \cap K$, then use the same procedure as above. Thus, the total number of vertices in the remaining graph is at most $|X| + 2k + 4k^3 = O(k^3)$. Furthermore, observe that adding and removing the four vertices as well as applying Reduction Rule 4.3 can be done in linear time. Thus, the overall running time of the kernelization is $O(n + m)$. □

4.2 Parameter feedback edge number

In this section, we provide a linear-time computable linear-size kernel of linear size for Matching parameterized by the feedback edge number, that is, the size of a feedback edge set. Observe that a feedback edge set can be computed in linear time via a simple depth-first search or breadth-first search. The kernel is based on the next two simple data reduction rules due to Karp and Sipser [19] and which deal with vertices of degree at most two.

Reduction Rule 4.5. *Let $v \in V$. If $\deg(v) = 0$, then delete v. If $\deg(v) = 1$, then delete v and its neighbor and decrease s by one (effectively, v is matched with its neighbor).*

Reduction Rule 4.6. *Let v be a vertex of degree two and let u, w be its neighbors. Then remove v, merge u and w, and decrease s by one.*

The correctness was stated by Karp and Sipser [19]. For completeness, we give a proof.

Lemma 21. *Reduction Rules 4.5 and 4.6 are correct.*

Proof. If v has degree zero, then clearly v cannot be in any matching and we can remove v.

If v has degree one, then let u be its single neighbor. Let M be a maximum matching of size at least s for G. Then v is matched in M since otherwise adding the edge $\{u, v\}$ would increase the size of the matching. Thus, a maximum matching in $G' = G - u - v$ is of size at least $s - 1$. Conversely, a maximum matching of size $s - 1$ in G' can easily be extended by the edge $\{u, v\}$ to a maximum matching of size s in G.

If v has degree two, then let u and w be its two neighbors. Let M a maximum matching of size at least s. If v is not matched in M, then u and w are matched since otherwise adding the edge $\{u, v\}$ resp. $\{v, w\}$ would increase the size of the matching. Thus, deleting v and merging u and w decreases the size of M by one (M looses either the edge incident to v or one of the edges incident to u and w). Hence, the resulting graph G' has a maximum matching of size at least $s - 1$. Conversely, let M' be a matching of size at least $s - 1$ for G'. If the

merged vertex vw is free, then $M := M' \cup \{\{u, v\}\}$ is a matching of size s in G. Otherwise, vw is matched to some vertex y. Then matching y in G with either v or w (at least one of the two vertices is a neighbor of y) and matching u with the other vertex yields a matching of size s for G. □

Although Reduction Rules 4.5 and 4.6 are correct, we do not know whether Reduction Rule 4.6 can be exhaustively applied in linear time. However, for our purpose it suffices to consider the following restricted version which we can exhaustively apply in linear time.

Reduction Rule 4.7. *Let v be a vertex of degree two and u, w its neighbors with u and w having degree at most two. Then remove v, merge u and w, and decrease s by one.*

Lemma 22. *Reduction Rules 4.5 and 4.7 can be exhaustively applied in $O(n + m)$ time.*

Proof. We give an algorithm which exhaustively applies Reduction Rules 4.5 and 4.7 in linear time. First, using bucket sort, sort the vertices by degree and keep three lists containing all degree-zero/one/two vertices. Then one applies Reduction Rules 4.5 and 4.7 in a straightforward way. When a neighbor of a vertex is deleted, then check if the vertex has now degree zero, one, or two. If yes, then add the vertex to the corresponding list.

We next show that this algorithm runs in linear time. First, observe that the deletion of each each degree-zero vertex can be done in constant time as no further vertices are affected. Second, consider a degree-one vertex v with a neighbor u and observe that deleting u and v can be done $O(\deg(v))$ time since one needs to update the degrees of all neighbors of v. Furthermore, decreasing s by one can be done in constant time for each deleted degree-one vertex. Finally, consider a degree-two vertex v with two neighbors u and w, each of degree at most two. Deleting v takes constant time. To merge u and w iterate over all neighbors of u and add them to the neighborhood of w. If a neighbor u' of u is already a neighbor of w, then decrease the degree of u' by one. Then, relabel w to be the new contracted vertex uw.

Overall, the worst-case running time to apply Reduction Rules 4.5 and 4.7 exhaustively can be upper-bounded by $O(n + \sum_{v \in V} \deg(v)) = O(n + m)$. □

Theorem 23. Matching *admits a linear-time computable linear-size kernel with respect to the parameter "feedback edge number".*

Proof. Apply Reduction Rules 4.5 and 4.7 exhaustively in linear time (see Lemma 22). We claim that the reduced graph $G = (V, E)$ has less than $12k$ vertices and $13k$ edges. Denote with $X \subseteq E$ a feedback edge set for G, $|X| \leq k$. Furthermore, denote with V^1_{G-X}, V^2_{G-X}, and $V^{\geq 3}_{G-X}$ the vertices that have degree one, two, and more than two in the $G - X$. Thus, $|V^1_{G-X}| \leq 2k$ as each leaf in $G - X$ has to be incident to an edge in X. Next, since $G - X$ is a forest (or tree), we have $|V^{\geq 3}_{G-X}| < |V^1_{G-X}|$ and thus $|V^{\geq 3}_{G-X}| < 2k$. Finally, each degree-two vertex in G needs at least one neighbor of degree at least three since G is reduced with respect to Reduction Rule 4.7. Thus, the vertices in V^2_{G-X} are either incident to an edge in X or adjacent to one of the at most $|V^{\geq 3}_{G-X}| + 2k$ vertices in G that have degree at least three. Since the sum over all degrees of vertices in $V^{\geq 3}_{G-X}$ is at most $\sum_{v \in V^{\geq 3}_{G-X}} \deg_{G-X}(v) \leq 2|V^{\geq 3}_{G-X}| + |V^1_{G-X}| < 6k$, it follows that $|V^2_{G-X}| \leq 8k$. Thus, the number of vertices in G is $|V^1_{G-X}| + |V^2_{G-X}| + |V^{\geq 3}_{G-X}| \leq 12k$. Since $G - X$ is a forest, it follows that G has at most $|V| + k \leq 13k$ edges. □

4.3 Parameter vertex cover number

We show a simple, quadratic-size kernel for MATCHING parameterized by the vertex cover number. Recall that a vertex cover is a vertex subset X such that each edge contains at least one vertex of X. Computing a factor-2 approximate vertex cover in linear time is folklore: As long as there is an edge in the graph add both endpoints into the vertex cover and delete both endpoints. Our kernelization uses Reduction Rule 4.5 and the following rule.

Reduction Rule 4.8. *Let $G = (V, E)$ be a graph and let $X \subseteq V$ be a vertex cover for G of size $|X| = k$. If a vertex $v \in X$ has more than k neighbors in $V \setminus X$, then delete the edges to all but k of these neighbors.*

Lemma 24. *Reduction Rule 4.8 is correct and can be exhaustively applied in linear time.*

Proof. We start with showing the correctness. Let $G = (V, E)$ be the input graph with the vertex cover X, $|X| = k$. Let $v \in X$ be a vertex in the vertex cover and let $v_1, \ldots, v_\ell$, $\ell > k$, be its neighbors in the independent set $V \setminus X$. The data reduction rule deletes the following set of edges $E' = \{\{v, v_{k+1}\}, \{v, v_{k+2}\}, \ldots, \{v, v_\ell\}\}$. We show that that there exists a maximum matching in G that does not contain any edge of E'. To this end, let M be a maximum matching for G that uses an edge from E', say $\{v, v_j\}$. Since the vertices in $V \setminus X$ form an independent set, it follows that they can only be matched to vertices in X. Since v is matched to v_j in M, it follows that at least one of the vertices $v_1, \ldots, v_k$, say v_i, has to be free with respect to M. Thus, $M' := (M \setminus \{\{v, v_j\}\}) \cup \{\{v, v_i\}\}$ is a maximum matching for G that uses no edge in E'. Hence, a maximum matching in the reduced instance has the same size as a maximum matching in the reduced graph.

The running time can be shown with a straightforward algorithm: For each vertex $v \in X$ iterate over all its neighbors and keep from its neighbors in $V \setminus X$ only the first k. Thus, each edge is visited once, resulting in $O(n + m)$ time. □

Theorem 25. MATCHING *parameterized by the vertex cover number k admits a linear-time computable kernel of size $O(k^2)$*

Proof. The kernelization is as follows: First exhaustively apply Reduction Rule 4.8 and then delete all isolated vertices (apply Reduction Rule 4.5). By Lemmas 22 and 24, this can be done in linear time.

It remains to show the kernel size: Since the kernel is reduced with respect to Reduction Rule 4.8, it follows that each vertex in the vertex cover X has at most k neighbors in $V \setminus X$. In particular, this implies that each vertex in X has degree less than $2k$. Since X is a vertex cover, the number of edges in the kernel is less than $2k^2$. Furthermore, since all isolated vertices are removed, the number of vertices $k^2 + k$. □

The kernel in the above theorem is relatively easy to obtain. However, generalizing the polynomial kernel to a stronger parameterization as "feedback vertex number" seems to be difficult, see the graph G depicted in Fig. 3 and the caption.

5 Almost linear-time kernelization

In this section we provide for MATCHING a kernel of size $O(k2^k)$ computable in $O(kn)$ time where k is the "feedback vertex number". Note that using a known a linear-time factor 4-approximation [1], we can approximate feedback vertex set and use it in our kernelization

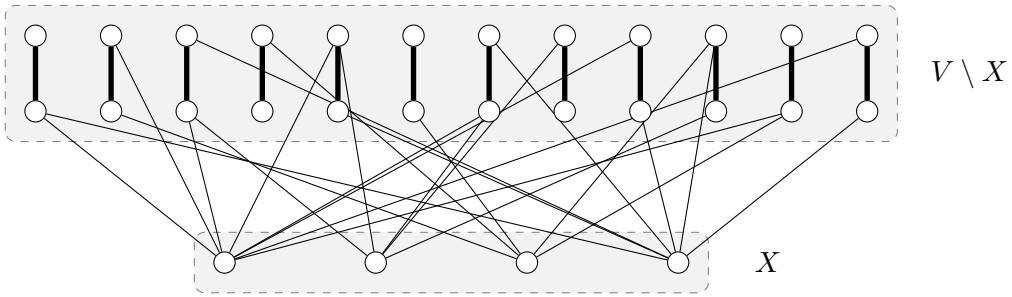

Figure 3: A difficult instance for kernelization with respect to the parameter $|X| = k$ (assume k to be divisible by two). The vertices in $V \setminus X$ above induce a perfect matching M. To keep the description simple, assume that the vertices in X are an independent set. The only augmenting paths in G are of the form u, u', v, v' where $u, v \in X$ and $\{u', v'\} \in M$. Extending the idea behind Reduction Rule 4.8 to kernelize this instance leads to the following: If for a pair $u, v \in X$ there are $k/2$ matched edges $\{u_1, v_1\}, \ldots, \{u_{k/2}, v_{k/2}\}$ such that for each of these edges u is adjacent to one endpoint u_i and v is adjacent to the other endpoint v_i, then delete all edges from u to vertices in $V \setminus \{u_1, \ldots, u_{k/2}\}$ and from v to vertices in $V \setminus \{v_1, \ldots, v_{k/2}\}$. Furthermore, if a vertex in $w \in V \setminus X$ has no neighbor in X, then delete w and its matched neighbor and reduce the solution size s by one. As each edge in M can be used in only one such augmenting path, keeping for each pair $u, v \in X$ at most $k/2$ such edges suffices: the pigeon hole principle ensures that at least one of these edges is "free" to be used by the pair u, v (and not used by another pair $u', v' \in X$). The problem with this kernelization approach is that it is too slow as a straightforward implementation requires $O(k^2 n)$ time and we can solve Matching in this time (see Theorem 4). In Section 5 we provide an exponential size kernel covering also the above shown instance.

algorithm. Furthermore, by Theorem 4, we can solve MATCHING in $O(k(n+m)) \leq O(k^2 n)$ time, since any graph with a feedback edge number k has $O(kn)$ edges.

Roughly speaking, our kernelization algorithm extends the linear-time computable kernel with respect to the parameter "feedback edge set". Thus, Reduction Rules 4.5 and 4.7 play an important role in the kernelization. Furthermore, the kernelization needs to overcome the problem depicted in Fig. 3. The solution presented in this section comes at the price of higher running time $O(kn)$ and bigger kernel size (exponential size) compared to the kernels presented in Section 4. It remains open whether MATCHING parameterized by the "feedback vertex number" admits a linear-time computable kernel (possibly of exponential size) or a polynomial kernel computable in $O(kn)$ time.

Subsequently, we describe our kernelization algorithm which keeps all vertices in the given feedback vertex set X and shrinks the size of $G - X$. Before doing so, we need some further notation. In this section, we assume that each tree is rooted at some arbitrary but fixed vertex so we can refer to the parent and children of a vertex. A leaf in $G - X$ is called a *bottommost leaf* if it has no siblings or all siblings are also leaves. (Here, bottommost refers to the subtree with the root being the parent of the considered leaf.) The outline of the algorithm is as follows (we assume throughout that $k < \log n$ since otherwise the input instance is already a kernel of size $O(2^k)$):

1. Reduce G wrt. Reduction Rules 4.5 and 4.7.
2. Compute a maximum matching M_{G-X} in $G - X$.
3. Modify M_{G-X} in linear time such that only the leaves of $G - X$ are free(Section 5.1).
4. Bound the number of free leaves in $G - X$ by k^2(Section 5.2).
5. Bound the number of bottommost leaves in $G - X$ by $O(k^2 2^k)$(Section 5.3).
6. Bound the degree of each vertex in X by $O(k^2 2^k)$. Then, use Reduction Rules 4.5 and 4.7 to provide the kernel of size $O(k^3 2^k)$(Section 5.4).

Whenever we reduce the graph in some step, we also show that the reduction is correct. That is, the given instance is a yes-instance if and only if the reduced one is a yes-instance. The correctness of our kernelization algorithm then follows by the correctness of each step. We discuss in the following the details of each step.

5.1 Steps 1 to 3

By Lemma 22 we can perform Step 1 in linear time. By Lemma 21 this step is correct.

Computing a maximum matching in Step 2 can be done by iteratively as follows: matching a free leaf to its neighbor and remove the two vertices from the graph. (Thus effectively applying Reduction Rule 4.5 to $G - X$.) By Lemma 22, this can be done in linear time.

Step 3 can be done in $O(n)$ time by traversing each tree in M_{G-X} in a BFS manner starting from the root: If a visited inner vertex v is free, then observe that all children are matched since M_{G-X} is maximum. Pick an arbitrary child u of v and match it with v. The vertex w that was previously matched to u is now free and since it is a child of u, it will be visited in the future. Observe that Steps 2 and 3 do not change the graph but only the auxiliary matching M_{G-X} and thus these steps are correct.

5.2 Step 4.

Recall that our goal is to upper-bound the number edges between vertices of X and $V \setminus X$ since we then can use the simple analysis used for the parameter "feedback edge set". Furthermore, recall that by Observation 2 the size of any maximum matching in G is at most k plus the size of M_{G-X}. Now, the crucial observation is that if a vertex $x \in X$ has at least k neighbors in $V \setminus X$ that are free wrt. M_{G-X}, then we know that there is maximum matching where x is matched to some one of these k vertices since at most $k-1$ can be "blocked" by other matching edges. (This is essentially the idea used in the kernel with respect to the parameter "vertex cover number", see Section 4.3.) This means we can delete all other edges incident to x. Formalizing this idea, we obtain the following data reduction rule (which is a generalization of Reduction Rule 4.8).

Reduction Rule 5.1. *Let $G = (V, E)$ be a graph, let $X \subseteq V$ be a subset of size k,and let M_{G-X} be a maximum matching for $G - X$. If there is a vertex $x \in X$ with at least k free neighbors $V_x = \{v_1, \ldots, v_k\} \subseteq V \setminus X$, then delete all edges from x to vertices in $V \setminus V_x$.*

Lemma 26. *Reduction Rule 5.1 is correct and can be exhaustively applied in $O(n+m)$ time.*

Proof. We first discuss the correctness and then the running time. Denote by s the size of a maximum matching in the input graph $G = (V, E)$ and by s' the size of a maximum matching in the new graph $G' = (V', E')$, where some edges incident to x are deleted. We need to show that $s = s'$. Since any matching in G' is also a matching in G, we easily obtain $s \geq s'$. It remains to show $s \leq s'$. To this end, let $M_G := M_G^{\max}(M_{G-X})$ be a maximum matching for G with the maximum overlap with M_{G-X} (see Section 2). If x is free wrt. M_G or if x matched to a vertex v that is also in G' a neighbor of x, then M_G is also a matching in G' ($M_G \subseteq E'$) and thus we have in this case $s \leq s'$. Hence, consider the remaining case where x is matched to some vertex v such that $\{v, x\} \notin E'$, that is, the edge $\{v, x\}$ was deleted by Reduction Rule 5.1. Hence, x has k neighbors $v_1, \ldots, v_k$ in $V \setminus X$ such that each of these neighbors is free wrt. M_{G-X} and none of the edges $\{v_i, x\}, i \in [k]$, was deleted. Observe that by the choice of M_G, the graph $G(M_{G-X}, M_G)$ (the graph over vertex set V and the edges that are either in M_{G-X} or in M_G, see Section 2) contains exactly $s - |M_{G-X}|$ paths (we do not consider isolated vertices as paths). Each of these paths is an augmenting path for M_{G-X}. By Observation 2, we have $s - |M_{G-X}| \leq k$. Observe that $\{v, x\}$ is an edge in one of these augmenting paths; denote this path with P. Thus, there are at most $k - 1$ paths $G(M_{G-X}, M_G)$ that do not contain x. Also, each of these paths contains exactly two vertices that are free wrt. M_{G-X}: the endpoints of the path. This means that no vertex in X is an inner vertex on such a path. Furthermore, since M_{G-X} is a maximum matching, it follows that for each path at most one of these two endpoints is in $V \setminus X$. Hence, at most $k-1$ vertices of $v_1, \ldots, v_k$ are contained in the $k-1$ paths of $G(M_{G-X}, M_G)$ except P. Therefore, one of these vertices, say v_i, is free wrt. M_G and can be matched with x. Thus, by reversing the augmentation along P and adding the edge $\{v_i, x\}$ we obtain another matching M'_G of size s. Observe that M'_G is a matching for G and for G' and thus we have $s \leq s'$. This completes the proof of correctness.

Now we come to the running time. We exhaustively apply the data reduction rule as follows. First, initialize for each vertex $x \in X$ a counter with zero. Second, iterate over all free vertices in $G - X$ in an arbitrary order. For each free vertex $v \in V \setminus X$ iterate over its neighbors in X. For each neighbor $x \in X$ do the following: if the counter is less than k,

then increase the counter by one and mark the edge $\{v, x\}$ (initially all edges are unmarked). Third, iterate over all vertices in X. If the counter of the currently considered vertex x is k, then delete all unmarked edges incident to x. This completes the algorithm. Clearly, it only deletes edges incident to a vertex $x \in X$ only if x has k free neighbors in $V \setminus X$ and the edges to these k neighbors are kept. The running time is $O(n + m)$: When iterating over all free vertices in $V \setminus X$ we consider each edge at most once. Furthermore, when iterating over the vertices in X, we again consider each edge at most once. □

To finish Step 4, we exhaustively apply Reduction Rule 5.1 in linear time. Afterwards, there are at most k^2 free (wrt. to M_{G-X}) leaves in $G-X$ that have at least one neighbor in X since each of the k vertices in X is adjacent to at most k free leaves. Thus, applying Reduction Rule 4.5 we can remove these free leaves. However, since for each degree-one vertex also its neighbor is removed, we might create new free leaves and need to again apply Reduction Rule 5.1 and update the matching (see Step 3). This process of alternating application of Reduction Rules 4.5 and 5.1 stops after at most k rounds since the neighborhood of each vertex in X can be changed by Reduction Rule 5.1 at most once. This shows the running time $O(k(n + m))$. We next show how to improve this to $O(n + m)$.

Lemma 27. *Given a matching instance (G, s) and a feedback vertex set X, Algorithm 3 computes in linear time an instance (G', s') with feedback vertex set X and a maximum matching $M_{G'-X}$ in $G' - X$ such that the following holds.*

- *There is a matching of size s in G if and only if there is a matching of size s' in G'.*
- *Each vertex that is free wrt. $M_{G'-X}$ is a leaf in $G' - X$.*
- *There are at most k^2 free leaves in $G' - X$.*

Proof. In the following, we explain Algorithm 3 which reduces the graph with respect to Reduction Rules 4.5 and 5.1 and updates the matching M_{G-X} as described in Step 3. The algorithm performs in (Algorithms 3 and 3) Steps 1 to 3. As described in the previous section, this can be done in linear time. Next, Reduction Rule 5.1 is applied in Algorithms 3 to 3 using the approach described in the proof of Lemma 26: For each vertex in x a counter $c(x)$ is maintained. When iterating over the free leaves in $G - X$, these counters will be updated. If a counter $c(x)$ reaches k, then the algorithm knows that x has k fixed free neighbors and according to Reduction Rule 5.1 the edges to all other vertices can be deleted (see Algorithm 3). Observe that once the counter $c(x)$ reaches k, the vertex x will never be considered again by the algorithm since its only remaining neighbors are free leaves in $G - X$ that already have been popped from stack L. The only difference from the description in the proof of Lemma 26 is that the algorithm reacts if the degree of some vertex v in $G - X$ is decreased to one (see Algorithms 3 to 3). If v is matched, then simply remove v and its matched neighbor from G and M_{G-X}. Otherwise, add v to the list L of unmatched degree-one vertices and defer dealing with v to a later stage of the algorithm.

Observe that the matching M_{G-X} still satisfies the property that each free vertex in $G-X$ is a leaf since only matched vertex pairs were deleted so far. When deleting unmatched degree-one vertices and their respective neighbor, the maximum matching M_{G-X} needs to be updated to satisfy this property. The algorithm does this from Algorithms 3 to 3: Let u be an entry in L such that u has degree one in Algorithm 3, that is, u is a free leaf in $G - X$ and

Algorithm 3: Reduce(G, M_{G-X}).

Input: A matching instance $(G = (V, E), s)$ and a feedback vertex set $X \subseteq V$ for G with $|X| = k$.

Output: An equivalent matching instance (G', s') such that X is also a feedback vertex set for G' and a maximum matching $M_{G'-X}$ for $G' - X$ such that only at most k^2 leaves in $G' - X$ are free.

```
1  Reduce G wrt. Reduction Rules 4.5 and 4.7
2  Compute a maximum matching M_{G-X} as described in Step 3
3  foreach x ∈ X do c(x) ← 0                    // c(x) will store the number of free neighbors for x
4  foreach e ∈ E do marked(e) ← False
5  L ← stack containing all free leaves in G − X
6  while L is not empty do
7      u ← pop(L)
8      foreach x ∈ N_G(u) ∩ X do                // Check whether Reduction Rule 5.1 is applicable for x
9          c(x) ← c(x) + 1, marked({u, x}) ← True          // fix u as free neighbor of x
10         if c(x) = k then                     // x has enough free neighbors: apply Reduction Rule 5.1
11             foreach y ∈ N_G(x) ∩ X do delete {x, y}
12             foreach v ∈ N_G(x) \ X do
13                 if marked({x, v}) = False then
14                     delete {x, v}
                       // Next deal with the case that v becomes a degree-one vertex
15                     if deg_G(v) = 1 and v is free then push v on L
16                     if deg_G(v) = 1 and v is matched then
17                         delete v and its neighbor, s ← s − 1
18     if deg_G(u) = 1 then                     // u has no neighbors in x
19         v ← neighbor of u in G − X; w ← matched neighbor of v
20         delete u and v from G
21         M_{G-X} ← M_{G-X} \ {{v, w}}, s ← s − 1          // update M_{G-X} and s
22         if w is now a leaf in G − X then
23             add w to L                       // w is free, so add w to the list of vertices to check
24         else
25             P ← arbitrary alternating path from w to a leaf u' in the subtree rooted in w
26             augment M_{G-X} along P          // ensure that the only free vertices are leaves
27             push u' on L                     // u' is free, add u' to the list of vertices to check
28 return (G, s) and M_{G-X}.
```

has no neighbors in X. Then, following Reduction Rule 4.5, delete u and its neighbor v and decrease the solution size s by one (see Algorithms 3 and 3). Let w denote the previously matched neighbor of v. Since v was removed, w is now free. If w is a leaf in $G - X$, then we can simply add it to L and in this way deal with it later. If w is not a leaf, then we need to update M_{G-X} since only leaves are allowed to be free. To this end, take an arbitrary alternating path P from w to a leaf u' of the subtree with root w and augment along P (see Algorithms 3 and 3). This can be done as follows: Pick an arbitrary child w_1 of w. Let w_2 be the matched neighbor of w_1. Since w is the parent of w_1, it follows that w_2 is a child of w_1. Now, remove $\{w_1, w_2\}$ from M_{G-X} and add $\{w_1, w\}$. If w_2 is a leaf, then the alternating path P is found with $u' = w_2$ and augmented. Otherwise, repeat the above

procedure with w_2 taking the role of w. This completes the algorithm. Its correctness follows from the fact that it only deletes edges and vertices according to Reduction Rules 4.5 and 5.1.

It remains to show the running time of $O(n+m)$. To this end, we prove that the algorithm considers each edge in E only two times. First, consider the edges incident to a vertex $x \in X$. These edges will be inspected at most twice by the algorithm: Once, when it is marked (see Algorithm 3). The second time is when it is deleted. This bounds the running time in the first part (Algorithms 3 to 3).

Now consider the remaining edges within $G - X$. To this end, observe that the algorithm performs two actions on the edges: deleting the edges (Algorithm 3) and finding and augmenting along an alternating path (Algorithms 3 and 3). Clearly, after deleting an edge it will no longer be considered, so it remains to show that edge is part of at most one alternating path used in Algorithms 3 and 3. Assume toward a contradiction that the algorithm augments along an edge twice or more. From all the edges that are augmented twice or more let $e \in E$ be one that is closest to the root of the tree e is contained in, that is, there is no edge closer to a root. Let w_1 and u'_1 be the endpoints of the first augmenting path P_1 containing e and w_2 and u'_2 the endpoints of the second augmenting path P_2 containing e. Observe that for each augmenting path chosen in Algorithm 3 it holds that one endpoint is a leaf and the other endpoint is an ancestor of this leaf. Assume without loss of generality that u'_1 and u'_2 are the leaves and w_1 and w_2 are their respective ancestors. Let u_1 and v_1 (u_2 and v_2) be the vertices deleted in Algorithm 3 which in turn made w_1 (w_2) free. Observe that e does not contain any of these four vertices u_1, v_1, u_2, v_2 since before augmenting P_1 (P_2) the vertices u_1 and v_1 (u_2 and v_2) are deleted. Since e is contained in both paths, either w_1 is an ancestor of w_2 or vice versa: the case $w_1 = w_2$ cannot happen since for the second augmenting path the endpoint $w_2 = w_1$ would not be matched to v_2; a contradiction (see Algorithm 3).

We next consider the case that w_1 is an ancestor of w_2 (the other case will be handled subsequently). Denote with w'_2 the neighbor of w_2 on P'. Observe that $e = \{w_2, w'_2\}$ since e is chosen as being closest to the root. We next distinguish the two cases whether or not e is initially matched. If e is initially free, then e is matched after augmenting along P_1. Then, by choice of P_2, e is not changed until the augmentation along P_2. This, however, is a contradiction since augmenting along P_2 only happens after the matched edge $\{w_2, v_2\}$ is deleted. Since $w'_2 \neq v_2$ and e is matched all the time until u_2 and v_2 are deleted, this means that w_2 would be matched to two vertices. Thus, consider the case that e is initially matched. Then, after augmenting along P_1, e is free and w_2 is matched to its parent p_{w_2}. As a consequence, v_2 is not matched to w_2, neither before nor after the augmentation of P_1. Since the algorithm augments along P_2 only after it deleted u_2 and v_2 where v_2 is matched to w_2, it follows that the edge $\{w_2, v_2\}$ is augmented before the algorithm augments along P_2. Denote with P_3 the augmenting path containing the edge $\{w_2, v_2\}$. Since w_2 is apparently not a free leaf, it follows that P_3 needs to contain the matched neighbor of w_2, which is p_{w_2}. This means that the edge $\{w_2, p_{w_2}\}$ is augmented at least twice (through P_1 and P_3). However, $\{w_2, p_{w_2}\}$ is closer to the root than $e = \{w_2, w'_2\}$, a contradiction to the choice of e. This completes the case that w_1 is an ancestor of w_2.

We now consider the remaining case where w_2 is an ancestor of w_1. In this case we have $e = \{w_1, w'_1\}$ where w'_1 is the neighbor of w_1 on P_1. Observe that w'_1 is a child of w_1. Furthermore, observe that after the augmentation along P_1 the leaf u'_1 is free and can be reached by an alternating path from w_1. Hence, before and after the augmentation along P_1 it holds that w_1 can reach exactly one free leaf via an alternating path (u_1 and u'_1). Observe that this is true even if the algorithm removes u'_1 since then a new free leaf will be created. Thus,

before deleting u_2 and v_2 (right before the augmentation along P_2), there is an augmenting path in $G - X$ from u_2 to w_1 and to the free leaf reachable from w_1. This is a contradiction to the fact that the matching M_{G-X} is maximum.

We conclude that each edge in E will be augmented at most once. Thus, the algorithm considers each edge at most twice (when augmenting it and when deleting it). Hence, the algorithm runs in linear time. □

Summarizing, in Step 4 we apply Algorithm 3 in order to obtain an instance with at most k^2 free vertices in $G - X$ that are all leaves. By Lemma 27 this can be done in linear time. Furthermore, Lemma 27 also shows that the step is correct.

5.3 Step 5

In this step we reduce the graph in $O(kn)$ time so that at most $k^2(2^k + 1)$ bottommost leaves will remain in the forest $G - X$. We will restrict ourselves to consider leaves that are matched with their parent vertex in M_{G-X} and that do not have a sibling. We call these bottommost leaves *interesting*. If a bottommost leaf has one or more siblings, then these sibling are by definition also leaves. Thus, at most one of these leaves (the bottommost leaf or its siblings) is matched with respect to M_{G-X} and the rest are free. Recall that in the previous step we upper-bounded the number of free leaves with respect to M_{G-X} by k^2 and thus there are at most k^2 bottommost leaves that are not interesting.

Our general strategy for this step is to extend the idea behind Reduction Rule 5.1: As described in Fig. 3, we want to keep for each pair of vertices $x, y \in X$ at most k different internally vertex-disjoint augmenting paths from x to y. (For ease of notation we keep k paths although keeping $k/2$ is sufficient.) In this step, we only consider augmenting paths of the form x, u, v, y where v is a bottommost leaf and u is v's parent in $G - X$. Assume that the parent u of v is adjacent to some vertex $x \in X$. Observe that in this case any augmenting path starting with the two vertices x and u has to continue to v and end in a neighbor of v. Thus, the edge $\{x, u\}$ can be only used in augmenting paths of length three. Furthermore, all these length-three augmenting paths are clearly internally vertex-disjoint. If we do not need the edge $\{x, u\}$ because we kept k augmenting paths from x already, then we can delete $\{x, u\}$. Furthermore, if we deleted the last edge from u to X (or u had no neighbors in X in the beginning), then u is a degree-two vertex in G and can be removed by applying Reduction Rule 4.6. As the child v of u is a leaf in $G - X$, it follows that v has at most $k + 1$ neighbors in G. We show below (Lemma 28) that the application of Reduction Rule 4.6 to remove u takes $O(k)$ time. As we remove at most n vertices, at most $O(kn)$ time is spent on Reduction Rule 4.6 in this step.

We now show that after a simple preprocessing one application of Reduction Rule 4.6 in the algorithm above can indeed be performed in $O(k)$ time.

Lemma 28. *Let u be a leaf in the tree $G - X$, v be its parent, and let w be the parent of v. If v has degree two in G, then applying Reduction Rule 4.6 to v (deleting v, contracting u and v, and setting $s := s - 1$) can be done in $O(k)$ time plus $O(kn)$ time for an initial preprocessing.*

Proof. The preprocessing is to simply create a partial adjacency matrix for G with the vertices in X in one dimension and V in the other dimension. This adjacency matrix has size $O(kn)$ and can clearly be computed in $O(kn)$ time.

Algorithm 4: FVS-Kernel Step 5.

Input: A matching instance $(G = (V, E), s)$, a feedback vertex set $X \subseteq V$ of size k for G with $k < \log n$, and a maximum matching M_{G-X} for $G - X$ with at most k^2 free vertices in $G - X$ that are all leaves.

Output: An equivalent matching instance (G', s') such that X is also a feedback vertex set for G' and $G - X$ is a tree with at most $k^2(2^k + 1)$ bottommost leaves, and a maximum matching $M_{G'-X}$ for $G' - X$ with at most k^2 free vertices in $G' - X$ that are all leaves.

```
1  Fix an arbitrary bijection f: 2^X → {1, ..., 2^k}
2  foreach v ∈ V \ X do
3      Set f_X(v) ← f(N(v) ∩ X)          // The number f_X(v) < n can be read in constant time.
4  Initialize a table Tab of size k · 2^k with Tab[x, f(Y)] ← 0 for all x ∈ X, ∅ ⊊ Y ⊆ X
5  P ← List containing all parents of interesting bottommost leaves
6  while P is not empty do
7      u ← pop(P)
8      v ← child vertex of u in G − X
9      foreach x ∈ N(u) ∩ X do
10         if Tab[x, f_X(v)] < k then
11             Tab[x, f_X(v)] ← Tab[x, f_X(v)] + 1
12         else
13             delete {x, u}
14     if u has now degree two in G then
15         Apply Reduction Rule 4.6 to u          // This decreases s by one.
16         vw ← vertex resulting from the merge of v and the parent w of u
17         if vw is now an interesting bottommost leaf then
18             add the parent of vw to P
19 return (G, s) and M_{G−X}.
```

Now apply Reduction Rule 4.6 to v. Deleting v takes constant time. To merge u and w iterate over all neighbors of u. If a neighbor u' of u is already a neighbor of w, then decrease the degree of u' by one, otherwise add u' to the neighborhood of w. Then, relabel w to be the new merged vertex uw.

Since u is a leaf in $G - X$ and its only neighbor in $G - X$, namely v, is deleted, it follows that all remaining neighbors of u are in X. Thus, using the above adjacency matrix, one can check in constant time whether u' is a neighbor of w. Hence, the above algorithm runs in $O(\deg(u)) = O(k)$ time. □

The above ideas are used in Algorithm 4 which we use for this step (Step 5). The algorithm is explained in the proof of the following lemma stating the correctness and the running time of Algorithm 4.

Lemma 29. *Let $(G = (V, E), s)$ be a matching instance, let $X \subseteq V$ be a feedback vertex set, and let M_{G-X} be a maximum matching for $G - X$ with at most k^2 free vertices in $G - X$ that are all leaves. Then, Algorithm 4 computes in $O(kn)$ time an instance (G', s') with feedback vertex set X and a maximum matching $M_{G'-X}$ in $G' - X$ such that the following holds.*

- *There is a matching of size s in G if and only if there is a matching of size s' in G'.*

- *There are at most $k^2(2^k + 1)$ bottommost leaves in $G' - X$.*
- *There are at most k^2 free vertices in $G' - X$ and they are all leaves.*

Proof. We start with describing the basic idea of the algorithm. To this end, let $\{u, v\} \in E$ be an edge such that v is an interesting bottommost leaf, that is, without siblings and matched to its parent u by M_{G-X}. Counting for each pair $x \in N(u) \cap X$ and $y \in N(v) \cap X$ one augmenting path gives in a simple worst-case analysis $O(k^2)$ time per edge, which is too slow for our purposes. Instead, we count for each pair consisting of a vertex $x \in N(u) \cap X$ and a set $Y = N(v) \cap X$ one augmenting path. In this way, we know that for each $y \in Y$ there is one augmenting path from x to y without iterating through all $y \in Y$. This comes at the price of considering up to $k2^k$ such pairs. However, we will show that we can do the computations in $O(k)$ time per considered edge in $G - X$ (instead of the $O(k^2)$ discussed in Fig. 3). The main reason for this improved running time is a simple preprocessing that allows for a bottommost vertex v to determine $N(v) \cap X$ in constant time.

The preprocessing is as follows (see Algorithms 4 to 4): First, fix an arbitrary bijection f between the set of all subsets of X to the numbers $\{1, 2, \ldots, 2^k\}$. This can be done for example by representing a set $Y \subseteq X = \{x_1, \ldots, x_k\}$ by a length-k binary string (a number) where the i^{th} position is 1 if and only if $x_i \in Y$. Given a set $Y \subseteq X$ such a number can be computed in $O(k)$ time in a straightforward way. Thus, Algorithms 4 to 4 can be performed in $O(kn)$ time. Furthermore, since we assume that $k < \log n$ (otherwise the input instance is already an exponential kernel), we have that $f(Y) < n$ for each $Y \subseteq X$. Thus, reading and comparing these numbers can be done in constant time. Furthermore, in Algorithm 4 the algorithm precomputes for each vertex the number corresponding to its neighborhood in X.

After the preprocessing, the algorithm uses a table Tab where it counts an augmenting path from a vertex $x \in X$ to a set $Y \subseteq X$ whenever a bottommost leaf v has exactly Y as neighborhood in X and the parent of v is adjacent to x (see Algorithms 4 to 4). To do this in $O(kn)$ time, the algorithm proceeds as follows: First, it computes in Algorithm 4 the set P which contains all parents of interesting bottommost leaves. Clearly, this can be done in linear time. Next, the algorithm processes the vertices in P. Observe that further vertices might be added to P (see Algorithm 4) during this processing. Let u be the currently processed vertex of P, let v be its child vertex, and let Y be the neighborhood of v in X. For each neighbor $x \in N(u) \cap X$, the algorithm checks whether there are already k augmenting paths between x and Y with a table lookup in Tab (see Algorithm 4). If not, then the table entry is incremented by one (see Algorithm 4) since u and v provide another augmenting path. If yes, then the edge $\{x, u\}$ is deleted in Algorithm 4 (we show below that this does not change the maximum matching size). If u has degree two after processing all neighbors of u in X, then by applying Reduction Rule 4.6, we can remove u and contract its two neighbors v and w. It follows from Lemma 28 that this application of Reduction Rule 4.6 can be done in $O(k)$ time. Hence, Algorithm 4 runs in $O(kn)$ time.

Recall that all vertices in $G - X$ that are free wrt. M_{G-X} are leaves. Thus, the changes to M_{G-X} by applying Reduction Rule 4.6 in Algorithm 4 are as follows: First, the edge $\{u, v\}$ is removed and second the edge $\{w, q\}$ is replaced by $\{vw, q\}$ for some $q \in V$. Hence, the matching M_{G-X} after running Algorithm 4 has still at most k^2 free vertices and all of them are leaves.

It remains to prove that (a) the deletion of the edge $\{x, u\}$ in Algorithm 4 results in an equivalent instance and (b) that the resulting instance has at most $O(k2^k)$ bottommost leaves. First, we show (a). To this end, assume towards a contradiction that the new graph $G' :=$

$G-\{x,u\}$ has a smaller maximum matching than G (clearly, G' cannot have a larger maximum matching). Thus, any maximum matching M_G for G has to contain the edge $\{x,u\}$. This implies that the child v of u in $G-X$ is matched in M_G with one of its neighbors (except u): If v is free wrt. M_G, then deleting $\{x,u\}$ from M_G and adding $\{v,u\}$ yields another maximum matching not containing $\{x,u\}$, a contradiction. Recall that $N(v) = \{u\} \cup Y$ where $Y \subseteq X$ since v is a leaf in $G-X$. Thus, each maximum matching M_G for G contains for some $y \in Y$ the edge $\{v,y\}$. Observe that Algorithm 4 deletes $\{x,u\}$ only if there are at least k other interesting bottommost leaves $v_1,\ldots,v_k$ in $G-X$ such that their respective parent is adjacent to x and $N(v_i) \cap X = Y$ (see Algorithms 4 to 4). Since $|Y| \leq k$, it follows by the pigeon hole principle that at least one of these vertices, say v_i, is not matched to any vertex in Y. Thus, since v_i is an interesting bottommost leaf, it is matched to its only remaining neighbor: its parent u_i in $G-X$. This implies that there is another maximum matching

$$M'_G := (M_G \setminus \{\{v,y\},\{x,u\},\{u_i,v_i\}\}) \cup \{\{v_i,y\},\{x,u_i\},\{u,v\}\},$$

a contradiction to the assumption that all maximum matchings for G have to contain $\{x,u\}$.

We next show (b) that the resulting instance has at most $k^2(2^k+1)$ bottommost leaves. To this end, recall that there are at most k^2 bottommost leaves that are not interesting (see discussion at the beginning of this subsection). Hence, it remains to upper-bound the number of interesting bottommost leaves. Observe that each parent u of an interesting bottommost leaf has to be adjacent to a vertex in X since otherwise u would have been deleted in Algorithm 4. Furthermore, after running Algorithm 4, each vertex $x \in X$ is adjacent to at most $k2^k$ parents of interesting bottommost leaves (see Algorithms 4 to 4). Thus, the number of interesting bottommost leaves is at most $k^2 2^k$. Therefore the number of bottommost leaves is upper-bounded by $k^2(2^k+1)$. □

5.4 Step 6

In this subsection, we provide the final step of the kernelization algorithm. Recall that in the previous steps we have upper-bounded the number of bottommost leaves in $G-X$ by $k^2 2^k$ and we computed a maximum matching M_{G-X} for $G-X$ such that at most k^2 vertices are free wrt. M_{G-X} and all free vertices are leaves in $G-X$. Using this, we show next how to reduce G to a graph of size $O(k^3 2^k)$. To this end, we need some further notation. A leaf in $G-X$ that is not bottommost is called a *pendant*. We define T to be the *pendant-free tree (forest) of* $G-X$, that is, the tree (forest) obtained from $G-X$ by removing all pendants. From the next observation, we can derive that $G-X$ is not much larger than T. This allows us to restrict ourselves in the following on giving an upper bound on the size of T.

Observation 30. *Let $G-X$ be as described above with vertex set $V \setminus X$ and let T be the pendant-free tree (forest) of $G-X$ with vertex set V_T. Then, $|V \setminus X| \leq 2|V_T| + k^2$.*

Proof. Observe that $V \setminus X$ is the union of all pendants in $G-X$ and V_T. Thus, it suffices to show that $G-X$ contains at most $|V_T|+k^2$ pendants. To this end, recall that we have a maximum matching for $G-X$ with at most k^2 free leaves. Thus, there are at most k^2 leaves in $G-X$ that have a sibling which is also a leaf since from two leaves with the same parent at most one can be matched. Hence, all but at most k^2 pendants in $G-X$ have pairwise different parent vertices. Since all these parent vertices are in V_T, it follows that the number of pendants in $G-X$ is $|V_T|+k^2$. □

We use the following observation to provide an upper bound on the number of leaves of T.

Observation 31. *Let F be a forest, let F' be the pendant-free forest of F, and let B be the set of all bottommost leaves in F. Then, the set of leaves in F' is exactly B.*

Proof. First observe that each bottommost leaf of F is a leaf of F' since we only remove vertices to obtain F' from F. Thus, it remains to show that each leaf v in F' is a bottommost leaf in F.

We distinguish two cases of whether or not v is a leaf in F: First, assume that v is not a leaf in F. Thus, all of it child vertices have been removed. Since we only remove pendants to obtain F' from F and since each pendant is a leaf, it follows that v is in F the parent of one or more leaves $u_1, \ldots, u_\ell$. Thus, by definition, all these leaves $u_1, \ldots, u_\ell$ are bottommost leaves, a contradiction to the fact that they were deleted when creating F'.

Second, assume that v is a leaf in F. If v is a bottommost leaf, then we are done. Thus, assume that v is not a bottommost leaf and therefore a pendant. However, since we remove all pendants to obtain F' from F, it follows that v is not contained in F', a contradiction. □

From Observation 31 it follows that the set B of bottommost leaves in $G-X$ is exactly the set of leaves in T. In the previous step we reduced the graph such that $|B| \leq k^2(2^k+1)$. Thus, T has at most $k^2(2^k+1)$ vertices of degree one and, since T is a tree (a forest), T also has at most $k^2(2^k+1)$ vertices of degree at least three. Let V_T^2 be the vertices of degree two in T and let $V_T^{\neq 2}$ be the remaining vertices in T. From the above it follows that $|V_T^{\neq 2}| \leq 2k^2(2^k+1)$. Hence, it remains to bound the size of V_T^2. To this end, we will upper-bound the degree of each vertex in X by $O(k^2 2^k)$ and then use Reduction Rules 4.5 and 4.7. We will check for each edge $\{x, v\} \in E$ with $x \in X$ and $V \setminus X$ whether we "need" it. This check will use the idea from the previous subsection where each vertex in X needs to reach each subset $Y \in X$ at most k times via an augmenting path. Similarly as in the previous section, we want to keep "enough" of these augmenting paths. However, this time the augmenting paths might be long and different augmenting paths can overlap. To still use the basic approach, we use the following lemma stating that we can still somehow replace augmenting paths.

Lemma 32. *Let M_{G-X} be a maximum matching in the forest $G-X$. Let P_{uv} be an augmenting path for M_{G-X} in G from u to v. Let P_{wx}, P_{wy}, and P_{wz} be three internally vertex-disjoint augmenting paths from w to x, y, and z, respectively, such that P_{uv} intersects all of them. Then, there exist two vertex-disjoint augmenting paths with endpoints u, v, w, and one of the three vertices x, y, and z.*

Proof. Label the vertices in P_{uv} alternating as *odd* or *even* with respect to P_{uv} so that no two consecutive vertices have the same label, u is *odd*, and v is *even*. Analogously, label the vertices in P_{wx}, P_{wy}, and P_{wz} as odd and even with respect to P_{wx}, P_{wy}, and P_{wz} respectively so that w is always odd. Since all these paths are augmenting, it follows that each edge from an even vertex to its succeeding odd vertex is in the matching M_{G-X} and each edge from an odd vertex to its succeeding even vertex is not in the matching. Observe that P_{uv} intersects the other paths in at least two vertices each since every second edge must be an edge in M_{G-X}. Since $G-X$ is a forest and all vertices in X are free with respect to M_{G-X}, it follows that the intersection of two augmenting paths is connected and thus a path. Since P_{uv} intersects the three augmenting paths from w, it follows that at least two of these paths, say P_{wx} and P_{wy}, have a "fitting parity", that is, in the intersections of P_{uv} with P_{wx} and with P_{wy} the even vertices with respect to P_{uv} are either even or odd with respect to *both* P_{wx} and P_{wy}.

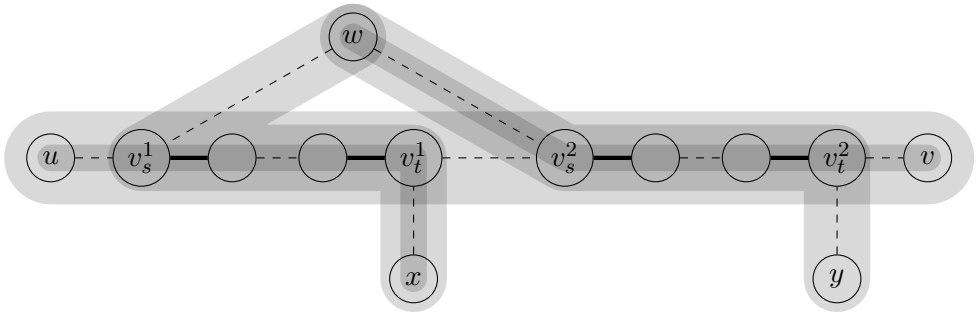

Figure 4: The situation in the proof of Lemma 32. The augmenting path from u to v intersects the two augmenting paths P_{wx} and P_{wy} from w to x and y respectively. Bold edges indicate edges in the matching, dashed edges indicate odd-length alternating paths starting with the first and last edge not being in the matching. The gray paths in the background highlight the different augmenting paths: the initial paths from u to v, w to x, and x to y as well as the new paths from u to x and w to v as postulated by Lemma 32.

Assume without loss of generality that in the intersections of the paths the vertices have the same label with respect to the three paths (if the labels differ, then revert the ordering of the vertices in P_{uv}, that is, exchange the names of u and v and change all labels on P_{uv} to its opposite). Denote with v_s^1 and v_t^1 the first and the last vertex in the intersection of P_{uv} and P_{wx}. Analogously, denote with v_s^2 and v_t^2 the first and the last vertex in the intersection of P_{uv} and P_{wy}. Assume without loss of generality that P_{uv} intersects first with P_{wx} and then with P_{wx}. Observe that v_s^1 and v_s^2 are even vertices and v_t^1 and v_t^2 are odd vertices since the intersections have to start and end with edges in M_{G-X} (see Fig. 4 for an illustration). Denote for a path P and two vertices p_1 and p_2 in P with $p_1 - P - p_2$ the subpath from p_1 to p_2. Observe that $u - P_{uv} - v_t^1 - P_{wx} - x$ and $w - P_{wy} - v_t^2 - P_{uv} - v$ are two vertex-disjoint augmenting paths. □

Algorithm description. A pseudocode is given in Algorithm 5, its details are at follows: The algorithm uses the same preprocessing (see Algorithms 4 to 4) as Algorithm 4. Thus, the algorithm can determine whether two vertices have the same neighborhood in X in constant time. As in Algorithm 4, Algorithm 5 uses a table Tab which has an entry for each vertex $x \in X$ and each set $Y \subseteq X$. The table is filled in such a way that the algorithm detected for each $y \in Y$ at least $\mathrm{Tab}[x, Y]$ internally vertex-disjoint augmenting paths from x to y.

The main part of the algorithm is the boolean function 'Keep-Edge' in Algorithms 5 to 5 which makes the decision on whether or not to delete an edge $\{x, v\}$ for $v \in V \setminus X$ and $x \in X$. The function works as follows for edge $\{x, v\}$: Starting at v the graph will be explored along possible augmenting paths until a "reason" for keeping the edge $\{x, v\}$ is found or until no further exploration is possible.

If the vertex v is free wrt. M_{G-X}, then $\{x, v\}$ is an augmenting path and we keep $\{x, v\}$ (see Algorithm 5). Observe that in Step 4 we upper-bounded the number of free vertices by k^2 and all these vertices are leaves. Thus, we keep a bounded number of edges incident to x because the corresponding augmenting paths can end in free leave. We provide the exact bound below when discussing the size of the graph returned by Algorithm 5. In Algorithm 5, the algorithm stops exploring the graph and keeps the edge $\{x, v\}$ if v has degree at least three

Algorithm 5: Algorithm for computing Step 6 of our kernel wrt. the parameter "feedback vertex number".

Input: A matching instance $(G = (V, E), s)$, a feedback vertex set $X \subseteq V$ of size k for G with $k < \log n$ and at most $k^2(2^k + 1)$ bottommost leaves in $G - X$, and a maximum matching M_{G-X} for $G - X$ with at most k^2 free vertices in $G - X$ that are all leaves.

Output: An equivalent matching instance (G', s') such that G' contains at most $O(k^3 2^k)$ vertices and edges.

Fix an arbitrary bijection $f \colon 2^X \to \{1, \ldots, 2^k\}$
foreach $v \in V \setminus X$ **do**
    Set $f_X(v) \leftarrow f(N(v) \cap X)$ // The number $f_X(v) < n$ can be read in constant time.
Initialize a table Tab of size $k \cdot 2^k$ with $\mathrm{Tab}[x, f(Y)] \leftarrow 0$ for $x \in X, \emptyset \subsetneq Y \subseteq X$
$T \leftarrow$ pendant-free tree (forest) of $G - X$
$V_T^{\geq 3} \leftarrow$ vertices in T with degree ≥ 3
foreach $x \in X$ **do**
    foreach $v \in N(x) \setminus X$ **do**
        if *Keep-Edge*$(x, v) =$ *false* **then** // Is $\{x, v\}$ needed for an augmenting path?
            delete $\{x, v\}$
exhaustively apply Reduction Rules 4.5 and 4.7
return (G, s).

Function *Keep-Edge*$(x \in X, v \in V \setminus X)$
    if *v is free wrt. M_{G-X} or $v \in V_T^{\geq 3}$* **then return** *true*
    $w \leftarrow$ matched neighbor of v in M_{G-X}
    if *$w \in V_T^{\geq 3}$ or w is adjacent to free leaf in $G - X$* **then return** *true*
    if *w has at least one neighbor in X and* $\mathrm{Tab}[x, f_X(w)] < 6k^2$ **then**
        $\mathrm{Tab}[x, f_X(w)] \leftarrow \mathrm{Tab}[x, f_X(w)] + 1$
        return *true*
    foreach *neighbor $u \neq v$ of w that is matched wrt. M_{G-X} and fulfills $\{u, x\} \notin E$* **do**
        if *Keep-Edge*$(u, x) =$ *true* **then return** *true*
    return *false*

in T. The reason for this is to keep the graph exploration simple in following only paths in T so that the running time for exploring the graph from x does not exceed $O(n)$. Furthermore, since the number of vertices in T with degree at least three is bounded (see discussion after Observation 31), it follows that only a bounded number of such edges $\{x, v\}$ are kept.

If v is not free wrt. M_{G-X}, then it is matched with some vertex w. If w is adjacent to some leaf u in $G - X$ that is free wrt. M_{G-X}, then the path x, v, w, u is an augmenting path. Thus, the algorithm keeps in this case the edge $\{x, v\}$, see Algorithm 5. Again, since the number of free leaves is bounded, only a bounded number of edges incident to x will be kept. If w has degree at least three in T, then the algorithm stops the graph exploration here and keeps the edge $\{x, v\}$, see Algorithm 5. Again, this is to keep the running time at $O(kn)$ overall.

Let $Y \subseteq X$ denote the neighborhood of w in X. Observe that the partial augmenting path x, v, w can be extended to each vertex in Y. Thus, if the algorithm did not yet find $6k^2$ paths from x to vertices whose neighborhood in X is also Y, then the table entry $\mathrm{Tab}[x, f_X(w)]$

(where $f_X(w)$ encodes the set $Y = N(w) \cap X$) is increased by one and the edge $\{x, v\}$ will be kept (see Algorithms 5 and 5). (Here we need $6k^2$ paths since these paths might be long and intersect with many other augmenting paths, see proof of Lemma 36 for the details of why $6k^2$ is enough.) If the algorithm already found $6k$ times an "augmenting path" from x to Y, then the neighborhood of w in X is not needed for x and the algorithm continues.

In Algorithm 5, all above discussed cases to keep the edge $\{x, v\}$ do not apply and the algorithm extends the partial augmenting part x, v, w by considering the neighbors of w except v. Since the algorithm dealt with possible extensions to vertices in X in Algorithms 5 to 5 and with extensions to free vertices in $G - X$ in Algorithm 5, it follows that the next vertex on this path has to be a vertex u that is matched wrt. M_{G-X}. Furthermore, since we want to extend a partial augmenting path from x, we require that u is not adjacent to x as otherwise x, u would be another, shorter partial augmenting path from x to u and we do not need the currently stored partial augmenting path.

Statements on Algorithm 5. For each edge $\{x, z\}$ with $x \in X$ and $z \in V \setminus X$ we denote by $P(x, z)$ the induced subgraph of $G - X$ on the vertices that are explored in the function Keep-Edge when called in Algorithm 5 with x and z. More precisely, we initialize $P(x, z) := \emptyset$. Whenever the algorithm reaches Algorithm 5, we add v to $P(x, z)$. Furthermore, whenever the algorithm reaches Algorithm 5, we add w to $P(x, z)$.

We next show that $P(x, z)$ is a path or a path with one additional pendant.

Lemma 33. *Let $x \in X$ and $z \in V \setminus X$ be two vertices such that $\{x, z\} \in E$. Then, $P(x, z)$ is either a path or a tree with exactly one vertex z' having more than two neighbors in $P(x, z)$. Furthermore, z' has degree exactly three and z is a neighbor of z'.*

Proof. We first show that all vertices in $P(x, z)$ except z and its neighbor z' have degree at most two in $P(x, z)$. Observe that having more vertices than z and z' in $P(x, z)$ requires Algorithm 5 to reach Algorithm 5.

Let w be the currently last vertex when Algorithm 5 continues the graph exploration in Algorithm 5. Observe that the algorithm therefore dealt with the case that w has degree at least three in T in Algorithm 5. Thus, w is either a pendant leaf in $G - X$ or $w \notin V_T^{\geq 3}$ (that is, w has degree at most two in T). In the first case, there is no candidate to continue and the graph exploration stops. In the second case, w has degree at most two in T.

We next show that any candidate u for continuing the graph exploration in Algorithm 5 is not a leaf in $G - X$. Assume toward a contradiction, that u is a leaf in $G - X$. Since the parent w of u is matched with some vertex $v \neq u$ (this is how w is chosen, see Algorithm 5), it follows that u is not matched. This implies that the function 'Keep-Edge' would have returned true in Algorithm 5 and would not have reached Algorithm 5, a contradiction. Thus, the graph exploration follows only vertices in T. Furthermore, the above argumentation implies that w is not adjacent to a leaf unless this leaf is its predecessor v in the graph exploration.

We now have two cases: Either w is not adjacent to a leaf in $G - X$ or $v = z$ is a leaf and $w = z'$ is its matched neighbor. In the first case, w has at most one neighbor $u \neq v$ since $w \notin V_T^{\geq 3}$. Hence, w has degree two in $P(x, z)$. In the second case, $w = z'$ has at most two neighbors $u \neq v$ and $u' \neq v$. Thus, z' has degree at most three. □

We set for $x \in X$

$$\mathcal{P}_x := \{P(x, v) \mid \{x, v\} \in E \wedge v \in V \setminus X\}$$

to be the union of all induced subgraphs wrt. x.

Lemma 34. *There exists a partition of $\mathcal{P}_x$ into $\mathcal{P}_x = \mathcal{P}_x^A \cup \mathcal{P}_x^B$ such that all graphs within $\mathcal{P}_x^A$ and within $\mathcal{P}_x^B$ are pairwise disjoint.*

Proof. Since $G - X$ is a tree (or forest), $G - X$ is also bipartite. Let A and B be its two color classes (so $A \cup B = V \setminus X$). We define the two parts $\mathcal{P}_x^A$ and $\mathcal{P}_x^B$ as follows: A subgraph $P \in \mathcal{P}_x$ is in $\mathcal{P}_x^A$ if the neighbor v of x in P is contained in A, otherwise P is in $\mathcal{P}_x^B$.

We show that all subgraphs in $\mathcal{P}_x^A$ and $\mathcal{P}_x^B$ are pairwise vertex-disjoint. To this end, assume toward a contradiction that two graphs $P, Q \in \mathcal{P}_x^A$ share some vertex. (The case $P, Q \in \mathcal{P}_x^B$ is completely analogous.) Let p_1 and q_1 be the first vertex in P and Q respectively, that is, p_1 and q_1 are adjacent to x in G. Observe that $p_1 \neq q_1$. Let $u \neq x$ be the first vertex that is in P and in Q. By Lemma 33, P and Q are paths or trees with at most one vertex of degree more than two and this vertex has degree three and is the neighbor of p_1 or q_1, respectively. This implies together with $q_1, p_1 \in A$ that either $u = p_1$ or $u = q_1$. Assume without loss of generality that $u = p_1$. Since $p_1 \in A$ and $q_1 \in A$ and u is a vertex in Q, it follows that Algorithm 5 followed u in the graph exploration from q_1 in Algorithm 5. However, this is a contradiction since the algorithm checks in Algorithm 5 whether the new vertex u in the path is not adjacent to x. Thus, all subgraphs in $\mathcal{P}_x^A$ and $\mathcal{P}_x^B$ are pairwise vertex-disjoint. □

We next show that if $\mathrm{Tab}[x, f(Y)] = 6k^2$ for some $x \in X$ and $Y \in X$ (recall that f maps Y to a number, see Algorithm 5), then there exist at least $3k$ internally vertex-disjoint augmenting paths from x to Y.

Lemma 35. *If in Algorithm 5 of Algorithm 5 it holds for $x \in X$ and $Y \subseteq X$ that $\mathrm{Tab}[x, f(Y)] = 6k^2$, then there exist in G wrt. M_{G-X} at least $3k$ alternating paths from x to vertices $v_1, \ldots, v_{3k^2}$ such that all these paths are pairwise vertex-disjoint (except x) and $N(v_i) \cap X = N(w) \cap X$ for all $i \in [3k^2]$.*

Proof. Note that each time $\mathrm{Tab}[x, f(Y)]$ is increased by one (see Algorithm 5), the algorithm found a vertex w such that there is an alternating path P from x to w and $N(w) \cap X = Y$. Furthermore, since the function Keep-Edge returns true in this case, the edge from x to its neighbor on P is not deleted in Algorithm 5. Thus, there exist at least $6k^2$ alternating paths from x to vertices whose neighborhood in X is exactly Y. By Lemma 34, it follows that at least half of these $6k^2$ paths are vertex-disjoint. □

The next lemma shows that Algorithm 5 is correct and runs in $O(kn)$ time.

Lemma 36. *Let $(G = (V, E), s)$ be a matching instance, let $X \subseteq V$ be a feedback vertex set of size k with $k < \log n$ and at most $k^2(2^k + 1)$ bottommost leaves in $G - X$, and let M_{G-X} be a maximum matching for $G - X$ with at most k^2 free vertices in $G - X$ that are all leaves. Then, Algorithm 5 computes in $O(kn)$ time an equivalent instance (G', s') of size $O(k^3 2^k)$.*

Proof. We split the proof into three claims, one for the correctness of the algorithm, one for the returned kernel size, and one for the running time.

Claim 37. *The input instance (G, s) is a yes-instance if and only if the instance (G', s') produced by Algorithm 5 is a yes-instance.*

Proof. Observe that the algorithm changes the input graph only in two lines: Algorithms 5 and 5. By Lemma 21, applying Reduction Rules 4.5 and 4.7 yields an equivalent instance. Thus, it remains to show that deleting the edges in Algorithm 5 is correct, that is, it does

not change the size of a maximum matching. To this end, observe that deleting edges does not increase the size of a maximum matching. Thus, we need to show that the size of the maximum matching does not decrease. Assume toward a contradiction that it does.

Let $\{x, v\}$ be the edge whose deletion decreased the maximum matching size. Redefine G to be the graph before the deletion of $\{x, v\}$ and G' to be the graph after the deletion of $\{x, v\}$. Recall that Algorithm 5 gets as additional input a maximum matching M_{G-X} for $G - X$. Let $M_G := M_G^{\max}(M_{G-X})$ be a maximum matching for G with the largest possible overlap with M and let $G^M := G(M_{G-X}, M_G) = (V, M_{G-X} \triangle M_G)$ (cf. Section 2). Since $\{x, v\} \in M_G \setminus M_{G-X}$ and x is free wrt. M_{G-X} it follows that there is a path P in G^M with one endpoint being x.

Recall that since P is a path in G^M it follows that P is an augmenting path for M_{G-X}. Since all vertices in X are free wrt. M_{G-X}, it follows that all vertices in P except the endpoints are in $V \setminus X$. Let z be the second endpoint of this path P. We call a vertex on P an even (odd) vertex if it has an even (odd) distance to x on P. (So x is an even vertex and v and z are odd vertices). Observe that v is the only odd vertex in P adjacent to x since otherwise there would be another augmenting path from x to z which only uses vertices from P implying the existence of another maximum matching that does not use $\{x, v\}$, a contradiction. Let u be the neighbor of z in P.

Since no odd vertex on P except v is adjacent to x, it follows that the graph exploration in the function Keep-Edge starting from x and v in Algorithm 5 either reached u or returned true before. If $z \in V \setminus X$, then in both cases, the function Keep-Edge would have returned true in Algorithm 5 and Algorithm 5 would not have deleted $\{x, v\}$, a contradiction. Thus, assume that $z \in X$. Therefore, the function Keep-Edge considered the vertex u in Algorithm 5 but did not keep the edge $\{x, v\}$. Thus, when considering u, it holds that $\mathrm{Tab}[x, f_X(u)] = 6k^2$, where $f_X(u)$ encodes $Y := N(u) \cap X$ and $z \in Y$.

By Lemma 35, it follows that there are $3k^2$ pairwise vertex-disjoint (except x) alternating paths from x to vertices $u_1, \ldots, u_k$ with $N(u_i) \cap X = Y$. Thus, there are $3k^2$ internally vertex-disjoint paths $\mathcal{Q}$ from x to y in G. If one of the paths $Q \in \mathcal{Q}$ does not intersect any path in G^M, then reverting the augmentation along P and augmenting along Q would results in another maximum matching not containing $\{x, v\}$, a contradiction. Thus, assume that each path in $\mathcal{Q}$ intersects at least one path in G^M.

For each two paths $Q_1, Q_2 \in \mathcal{Q}$ that are intersected by the same path P' in G^M it holds that each further path P'' in G^M can intersect at most one of Q_1 and Q_2: Assume toward a contradiction that P'' does. Since no path in G^M except P contains x and z it follows that all intersections between the paths are within $G - X$. Since P' and P'' are vertex-disjoint and Q_1 and Q_2 are internally vertex-disjoint, it follows that there is a cycle in $G - X$, a contradiction to the fact that X is a feedback vertex set.

Since $3k^2 > 3k + k^2$, it follows from the pigeon hole principle that there is a path $P' \in G^M$ that intersects at least three paths $Q_1, Q_2, Q_3 \in \mathcal{Q}$ such that no further path in G^M intersects them. We can now apply Lemma 32 and obtain two vertex-disjoint augmenting paths Q and Q'. Thus, reverting the augmentation along P and P'' and augment along Q and Q'' yields another maximum matching for G which does not contain $\{x, v\}$, a contradiction. □

Claim 38. *The graph G' returned by Algorithm 5 has at most $O(k^3 2^k)$ vertices and edges.*

Proof. We first show that each vertex $x \in X$ has degree at most $O(k^2 2^k)$ in G'. To this end, we need to count the number of neighbors $v \in N(x) \setminus X$ where the function Keep-Edge returns true in Algorithm 5. By Lemma 33, the function Keep-Edge explores the graph along one or

two paths (essentially growing from one starting point into two directions). Recall that $\mathcal{P}_x$ denotes the subgraphs induced by the graph exploration of Keep-Edge for the neighbors of x. By Lemma 34 there is a partition of $\mathcal{P}_x$ into $\mathcal{P}_x^A$ and $\mathcal{P}_x^B$ such that within each part the subgraphs are pairwise vertex-disjoint. We consider the two parts independently. We start with bounding the number of graphs in $\mathcal{P}_x^A$ where the function 'Keep-Edge' returned true (the analysis is completely analog for $\mathcal{P}_x^B$).

Since all explored subgraphs are disjoint and all free vertices in $G - X$ wrt. M_{G-X} are leaves, it follows that Algorithm 5 returned at most k^2 times true in Algorithm 5 due to w being adjacent to a free leaf in $G - X$. Also, the algorithm returns at most k^2 times true in Algorithm 5 due to v being free. Furthmore, the algorithm returns at most $6k^2 \cdot 2^k$ times true in Algorithm 5. Finally, we show that the algorithm returns at most $k^2 \cdot (2^k - 1)$ times true in Algorithms 5 and 5, respectively. It follows from the discussion below Observation 31 that T, the pendent-free tree of $G - X$, has at most $k^2(2^k + 1)$ leaves (denoted by V_T^1) and $k^2(2^k + 1)$ vertices of degree at least three (denoted by $V_T^{\geq 3}$). Let V_T be the vertices of T. Since T is a tree (or forest), it has more vertices than edges and hence

$$\sum_{v \in V_T} \deg_T(v) < 2|V_T|$$

which implies

$$\sum_{v \in V_T^{\geq 3}} \deg_T(v) < 2 \cdot |V_T^{\geq 3}| + |V_T^1|.$$

Thus, Algorithm 5 returns at most $2 \cdot |V_T^{\geq 3}| + |V_T^1| < 3k^2(2^k + 1)$ times true in Algorithm 5 due to w being a vertex in $V_T^{\geq 3}$. Also, Algorithm 5 returns at most $|V_T^{\geq 3}| \leq k^2(2^k + 1)$ times true in Algorithm 5 due to v being a vertex in $V_T^{\geq 3}$.

Summarizing, considering the graph explorations in $\mathcal{P}_x^A$, Algorithm 5 returned at most

$$k^2 + k^2 + 6k^2 \cdot 2^k + 2k^2(2^k + 1) \in O(k^2 2^k)$$

times true in the function Keep-Edge. Analogously, considering the graph explorations in $\mathcal{P}_x^A$, Algorithm 5 also returned at most $O(k^2 2^k)$ times true. Hence, each vertex $x \in X$ has degree at most $O(k^2 2^k)$ in G'.

We now show that the exhaustive application of Reduction Rules 4.5 and 4.7 indeed results in a kernel of the claimed size. To this end, denote with $V_{G'-X}^1$, $V_{G'-X}^2$, and $V_{G'-X}^{\geq 3}$ the vertices that have degree one, two, and at least three in $G' - X$. We have $|V_{G'-X}^1| \in O(k^3 2^k)$ since each vertex in X has degree at most $O(k^2 2^k)$ and G' is reduced wrt. Reduction Rule 4.5. Next, since $G' - X$ is a forest (or tree), we have $|V_{G'-X}^{\geq 3}| < |V_{G'-X}^1|$ and thus $|V_{G'-X}^{\geq 3}| \in O(k^3 2^k)$. Finally, each degree-two vertex in G' needs at least one neighbor of degree at least three since G' is reduced with respect to Reduction Rule 4.7. Thus, each vertex in $V_{G'-X}^2$ is either incident to a vertex in X or adjacent to one of the at most $O(k^2 2^k)$ vertices in $G' - X$ that have degree at least three. Thus, $|V_{G'-X}^2| \in O(k^3 2^k)$. Summarizing, G' contains at most $O(k^3 2^k)$ vertices and edges. □

Claim 39. *Algorithm 5 runs in $O(kn)$ time.*

Proof. First, observe that Algorithms 5 to 5 can be done in $O(kn)$ time: The preprocessing and table initialization can be done in $O(kn)$ time as discussed in Section 5.3. Furthermore, T and $V_T^{\geq 3}$ can clearly be computed in $O(n + m) \leq O(kn)$ time. Second, by Lemma 22,

applying Reduction Rules 4.5 and 4.7 can be done in $O(n+m)$ time. Thus, it remains to show that each iteration of the foreach-loop in Algorithm 5 can be done in $O(n)$ time.

By Lemma 34, the explored graphs $\mathcal{P}_x$ from x can be partitioned into two parts such that within each part all subgraphs are vertex-disjoint. Thus, each vertex in $G-X$ is visited only twice during the execution of the function Keep-Edge. Furthermore, observe that in Algorithms 5 and 5 the table can be accessed in constant time. Thus, the function Keep-Edge only checks once whether a vertex in $V \setminus X$ has a neighbor in X, namely in Algorithm 5. This single check can be done in constant time. Since the rest of the computation is done on $G-X$ which has less than $|V \setminus X|$ edges, it follows that each iteration of the foreach-loop in Algorithm 5 can indeed be done in $O(n)$ time. □

This completes the proof of Lemma 36. □

Our kernelization algorithm for the parameter "feedback vertex number" essentially calls Steps 1 to 6.

Theorem 40. Matching *parameterized by the* feedback vertex number k *admits a kernel of size* $2^{O(k)}$. *It can be computed in* $O(kn)$ *time.*

Proof. First, using the linear-time factor-four approximation of Bar-Yehuda et al. [1], we compute an approximate feedback vertex set X with $|X| \leq 4k$. Then, we apply Steps 1 to 6. Applying the first three steps is rather straightforward, see Section 5.1. For the remaining three steps, we use Algorithms 3 to 5. By Lemmas 27, 29 and 36, this can be done in $O(kn)$ time and results in a kernel of size $O((4k)^3 2^{4k}) = 2^{O(k)}$. □

6 Conclusion

There were two main targets of this work. On the one side, we tried to further explore the computational complexity landscape of finding maximum matchings. More specifically, the leitmotif was to find out under what circumstances we can beat the well-known time bound $(m\sqrt{n})$ by exact algorithms. On the other side, we tried to further chart the feasibilty of the "FPT in P" program recently initiated [15]. To this end, we attacked a central polynomial-time solvable problem and tried to see how well concepts of parameterized and multivariate algorithm design (in particular, kernelization) apply. The bottom line for both our targets is that we could gain new insights, new and prospective algorithms (their practical relevance is to be validated in more empirical follow-up work), and quite a number of challenging theoretical and combinatorial questions and directions.

We conclude with a few specific questions and tasks for future research. Relating to the described framework of gaining linear-time fixed-parameter algorithms based on distance parameterizations to graph classes, can we compute (or approximate) in linear time a minimum-size vertex subset X such that $G-X$ is a cocomparability graph? Moreover, is there a linear-time computable kernel for Matching parameterized by the distance to cocomparability graphs (assuming that the distance set is given)? Can we extend the kernel of Section 4.1 from Bipartite Matching to Matching parameterized by the distance to chain graphs? These are only three very concrete challenges. There are numerous more, including further parameterizations or any lower bound results in our multivariate setting.

References

[1] R. Bar-Yehuda, D. Geiger, J. Naor, and R. M. Roth. "Approximation algorithms for the feedback vertex set problem with applications to constraint satisfaction and Bayesian inference". In: *SIAM Journal on Computing* 27.4 (1998), pp. 942–959 (cit. on pp. 6, 24, 42).

[2] N. Blum. "A New Approach to Maximum Matching in General Graphs". In: *Proceedings of the 17th International Colloquium on Automata, Languages, and Programming (ICALP '90)*. Vol. 443. LNCS. Springer, 1990, pp. 586–597 (cit. on p. 5).

[3] A. Brandstädt, V. B. Le, and J. P. Spinrad. *Graph Classes: a Survey*. Vol. 3. SIAM Monographs on Discrete Mathematics and Applications. SIAM, 1999 (cit. on pp. 7, 14).

[4] L. Cai. "Parameterized Complexity of Vertex Colouring". In: *Discrete Applied Mathematics* 127.1 (2003), pp. 415–429 (cit. on pp. 3, 6).

[5] M. Chang. "Algorithms for Maximum Matching and Minimum Fill-in on Chordal Bipartite Graphs". In: *Proceedings of the 7th International Symposium on Algorithms and Computation (ISAAC '96)*. Vol. 1178. LNCS. Springer, 1996, pp. 146–155 (cit. on p. 2).

[6] D. G. Corneil, B. Dalton, and M. Habib. "Certifying algorithm for the Minimum Path Cover problem on cocomparability graphs using LDFS". In: *SIAM Journal on Computing* 42.3 (2013), 792—807 (cit. on pp. 4, 8, 9).

[7] D. G. Corneil and R. M. Krueger. "A Unified View of Graph Searching". In: *SIAM Journal on Discrete Mathematics* 22.4 (2008), pp. 1259–1276 (cit. on pp. 3, 8).

[8] E. Dahlhaus and M. Karpinski. "Matching and multidimensional matching in chordal and strongly chordal graphs". In: *Discrete Applied Mathematics* 84.1–3 (1998), pp. 79 –91 (cit. on p. 2).

[9] R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity*. Springer, 2013 (cit. on p. 13).

[10] R. Duan and S. Pettie. "Linear-Time Approximation for Maximum Weight Matching". In: *Journal of the ACM* 61.1 (2014), 1:1–1:23 (cit. on pp. 1, 2).

[11] F. V. Fomin, D. Lokshtanov, M. Pilipczuk, S. Saurabh, and M. Wrochna. "Fully polynomial-time parameterized computations for graphs and matrices of low treewidth". In: *CoRR* abs/1511.01379 (2015) (cit. on p. 2).

[12] H. N. Gabow and R. E. Tarjan. "A Linear-Time Algorithm for a Special Case of Disjoint Set Union". In: *Journal of Computer and System Sciences* 30.2 (1985), pp. 209–221 (cit. on p. 2).

[13] H. N. Gabow and R. E. Tarjan. "Faster Scaling Algorithms for General Graph-Matching Problems". In: *Journal of the ACM* 38.4 (1991), pp. 815–853 (cit. on p. 5).

[14] F. Gardi. "Efficient Algorithms for Disjoint Matchings among Intervals and Related Problems". In: *Proceedings of the 4th International Conference on Discrete Mathematics and Theoretical Computer Science (DMTCS '03)*. Vol. 2731. LNCS. Springer, 2003, pp. 168–180 (cit. on p. 2).

[15] A. C. Giannopoulou, G. B. Mertzios, and R. Niedermeier. "Polynomial Fixed-Parameter Algorithms: A Case Study for Longest Path on Interval Graphs". In: *CoRR* abs/1506.01652 (2015) (cit. on pp. 2, 42).

[16] A. V. Goldberg and A. V. Karzanov. “Maximum skew-symmetric flows and matchings”. In: *Mathematical Programming* 100.3 (2004), pp. 537–568 (cit. on p. 2).

[17] J. Guo, F. Hüffner, and R. Niedermeier. “A Structural View on Parameterizing Problems: Distance from Triviality”. In: *Proceedings of the 1st International Workshop on Parameterized and Exact Computation (IWPEC ’04)*. Vol. 3162. LNCS. Springer, 2004, pp. 162–173 (cit. on pp. 3, 6).

[18] J. E. Hopcroft and R. M. Karp. “An $n^{5/2}$ Algorithm for Maximum Matchings in Bipartite Graphs”. In: *SIAM J. Comput.* 2.4 (1973), pp. 225–231 (cit. on pp. 2, 5).

[19] R. M. Karp and M. Sipser. “Maximum Matchings in Sparse Random Graphs”. In: *Proceedings of the 22nd Annual IEEE Symposium on Foundations of Computer Science (FOCS ’81)*. IEEE Computer Society, 1981, pp. 364–375 (cit. on p. 22).

[20] K. Kaya, J. Langguth, F. Manne, and B. Uçar. *Experiments on push-relabel-based maximum cardinality matching algorithms for bipartite graphs*. Tech. rep. Technical Report TR/PA/11/33. CERFACS, Toulouse, France, 2011 (cit. on p. 2).

[21] E. Köhler and L. Mouatadid. “Linear time LexDFS on cocomparability graphs”. In: *Proceedings of the 14th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT)*. 2014, pp. 319–330 (cit. on p. 9).

[22] C. Komusiewicz and R. Niedermeier. “New Races in Parameterized Algorithmics”. In: *Proceedings of the 37th International Symposium on Mathematical Foundations of Computer Science (MFCS ’12)*. Vol. 7464. LNCS. Springer, 2012, pp. 19–30 (cit. on pp. 3, 6).

[23] D. Kratsch and L. Stewart. “Domination on cocomparability graphs”. In: *SIAM Journal on Discrete Mathematics* 6.3 (1993), pp. 400–417 (cit. on p. 8).

[24] Y. D. Liang and C. Rhee. “Finding a Maximum Matching in a Circular-Arc Graph”. In: *Inf. Process. Lett.* 45.4 (1993), pp. 185–190 (cit. on p. 2).

[25] A. Madry. “Navigating Central Path with Electrical Flows: From Flows to Matchings, and Back”. In: *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS ’13)*. IEEE Computer Society, 2013, pp. 253–262 (cit. on p. 2).

[26] R. M. McConnell and J. Spinrad. “Linear-time modular decomposition and efficient transitive orientation of comparability graphs”. In: *Proceedings of the tenth annual ACM-SIAM symposium on Discrete algorithms (SODA)*. 1994, pp. 536–545 (cit. on p. 8).

[27] G. B. Mertzios and D. G. Corneil. “A simple polynomial algorithm for the longest path problem on cocomparability graphs”. In: *SIAM Journal on Discrete Mathematics* 26.3 (2012), pp. 940–963 (cit. on pp. 4, 8).

[28] S. Micali and V. V. Vazirani. “An $O(\sqrt{|V|}|E|)$ Algorithm for Finding Maximum Matching in General Graphs”. In: *Proceedings of the 21st Annual IEEE Symposium on Foundations of Computer Science (FOCS ’80)*. IEEE, 1980, pp. 17–27 (cit. on pp. 1, 2, 5).

[29] M. Mucha and P. Sankowski. “Maximum Matchings in Planar Graphs via Gaussian Elimination”. In: *Algorithmica* 45.1 (2006), pp. 3–20 (cit. on p. 2).

[30] M. Mucha and P. Sankowski. “Maximum Matchings via Gaussian Elimination”. In: *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS '04)*. IEEE Computer Society, 2004, pp. 248–255 (cit. on p. 2).

[31] S. Olariu. “An Optimal Greedy Heuristic to Color Interval Graphs”. In: *Information Processing Letters* 37.1 (1991), pp. 21–25 (cit. on p. 7).

[32] P. Sankowski. “Maximum weight bipartite matching in matrix multiplication time”. In: *Theoretical Computer Science* 410.44 (2009), pp. 4480 –4488 (cit. on p. 2).

[33] S. S. Skiena. *The Algorithm Design Manual*. Springer, 2010 (cit. on p. 1).

[34] J. P. Spinrad. *Efficient graph representations*. Vol. 19. Fields Institute Monographs. American Mathematical Society, 2003 (cit. on p. 8).

[35] G. Steiner and J. S. Yeomans. “A linear time algorithm for maximum matchings in convex bipartite graphs”. In: *Comput. Math. Appl.* 31 (1996), pp. 91–96 (cit. on pp. 2, 5, 14, 15).

[36] M. Yu and C. Yang. “An $O(n)$ Time Algorithm for Maximum Matching on Cographs”. In: *Information Processing Letters* 47.2 (1993), pp. 89–93 (cit. on p. 2).

[37] R. Yuster. “Maximum Matching in Regular and Almost Regular Graphs”. In: *Algorithmica* 66.1 (2013), pp. 87–92 (cit. on p. 2).